NASA/CR-2000-210105

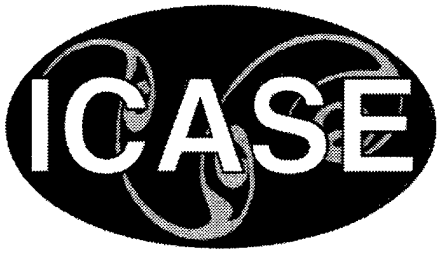

Semiannual Report

October 1, 1999 through March 31, 2000

DISTRIBUTION STATEMENT A
Approved for Public Release
Distribution Unlimited

April 2000

20000531 043

DTIC QUALITY INSPECTED 1

The NASA STI Program Office . . . in Profile

Since its founding, NASA has been dedicated to the advancement of aeronautics and space science. The NASA Scientific and Technical Information (STI) Program Office plays a key part in helping NASA maintain this important role.

The NASA STI Program Office is operated by Langley Research Center, the lead center for NASA's scientific and technical information. The NASA STI Program Office provides access to the NASA STI Database, the largest collection of aeronautical and space science STI in the world. The Program Office is also NASA's institutional mechanism for disseminating the results of its research and development activities. These results are published by NASA in the NASA STI Report Series, which includes the following report types:

- TECHNICAL PUBLICATION. Reports of completed research or a major significant phase of research that present the results of NASA programs and include extensive data or theoretical analysis. Includes compilations of significant scientific and technical data and information deemed to be of continuing reference value. NASA counter-part or peer-reviewed formal professional papers, but having less stringent limitations on manuscript length and extent of graphic presentations.

- TECHNICAL MEMORANDUM. Scientific and technical findings that are preliminary or of specialized interest, e.g., quick release reports, working papers, and bibliographies that contain minimal annotation. Does not contain extensive analysis.

- CONTRACTOR REPORT. Scientific and technical findings by NASA-sponsored contractors and grantees.

- CONFERENCE PUBLICATIONS. Collected papers from scientific and technical conferences, symposia, seminars, or other meetings sponsored or co-sponsored by NASA.

- SPECIAL PUBLICATION. Scientific, technical, or historical information from NASA programs, projects, and missions, often concerned with subjects having substantial public interest.

- TECHNICAL TRANSLATION. English-language translations of foreign scientific and technical material pertinent to NASA's mission.

Specialized services that help round out the STI Program Office's diverse offerings include creating custom thesauri, building customized databases, organizing and publishing research results . . . even providing videos.

For more information about the NASA STI Program Office, you can:

- Access the NASA STI Program Home Page at http://www.sti.nasa.gov/STI-homepage.html

- Email your question via the Internet to help@sti.nasa.gov

- Fax your question to the NASA Access Help Desk at (301) 621-0134

- Phone the NASA Access Help Desk at (301) 621-0390

- Write to:
  NASA Access Help Desk
  NASA Center for AeroSpace Information
  7121 Standard Drive
  Hanover, MD 21076-1320

Available from the following:

NASA Center for AeroSpace Information (CASI)
7121 Standard Drive
Hanover, MD 21076-1320
(301) 621-0390

National Technical Information Service (NTIS)
5285 Port Royal Road
Springfield, VA 22161-2171
(703) 487-4650

CONTENTS

Page

Introduction ii

Research in Progress

Applied and Numerical Mathematics 1

Physical Sciences 10

Computer Science 17

Reports and Abstracts 25

ICASE Colloquia 35

ICASE Staff 37

INTRODUCTION

The Institute for Computer Applications in Science and Engineering (ICASE)* is operated at the Langley Research Center (LaRC) of NASA by the Universities Space Research Association (USRA) under a contract with the Center. USRA is a nonprofit consortium of major U.S. colleges and universities.

The Institute conducts unclassified basic research in applied mathematics, numerical analysis and algorithm development, computer science, fluid mechanics, and structures and materials in order to extend and improve problem-solving capabilities in science and engineering, particularly in the areas of aeronautics and space research.

ICASE has a small permanent staff. Research is conducted primarily by visiting scientists from universities and industry who have resident appointments for limited periods of time as well as by visiting and resident consultants. Members of NASA's research staff may also be residents at ICASE for limited periods.

The major categories of the current ICASE research program are:

- Applied and numerical mathematics, including multidisciplinary design optimization;
- Applied computer science: system software, systems engineering, and parallel algorithms;
- Theoretical, computational, and experimental research in fluid mechanics in selected areas of interest to LaRC, such as transition, turbulence, flow control, and acoustics; and
- Theoretical, computational, and experimental research in structures and material sciences with emphasis on smart materials.

ICASE reports are primarily considered to be preprints of manuscripts that have been submitted to appropriate research journals or that are to appear in conference proceedings. A list of these reports for the period October 1, 1999 through March 31, 2000 is given in the Reports and Abstracts section, which follows a brief description of the research in progress.

*ICASE is operated at NASA Langley Research Center, Hampton, VA, under the National Aeronautics and Space Administration, NASA Contract No. NAS1-97046. Financial support was provided by NASA Contract Nos. NAS1-97046, NAS1-19480, NAS1-18605, NAS1-18107, NAS1-17070, NAS1-17130, NAS1-15810, NAS1-16394, NAS1-14101, and NAS1-14472.

RESEARCH IN PROGRESS

APPLIED AND NUMERICAL MATHEMATICS

Active Flow Control Research Tool Development

Brian Allan

The use of secondary flow control devices on inlets, such as vortex generators and small synthetic jets, has recently attracted interest. Significant improvements in engine flow distortion and pressure loss have been demonstrated using these small flow control devices. This technology has shown a great potential in improving existing inlet configurations and enabling more advanced inlet designs, which allow for reduced inlet lengths and sharper turning angles. In order to take advantage of these devices, computational tools need to be developed to evaluate potential engine inlet designs. The goal of this research will be to develop and test computational boundary conditions and techniques to model vortex generators and synthetic jets, for CFD codes, which are to be tested and evaluated on an inlet application.

Using a compressible Navier-Stokes code, developed at NASA, numerical simulations of micro-vortex generators and synthetic jets on a flat plate and inlet duct flows, will be performed. These simulations are to be performed in conjunction with wind tunnel experiments conducted at NASA Langley Research Center. From these simulations and experiments, a better understanding of the flow physics of the flow control devices will be gained. Using these insights into the flow physics of the control devices, a computational boundary condition can then be developed. This boundary condition will capture the general behavior of the flow control devices without having to simulate them directly. This will result in a reduced computational cost by eliminating the need to include fine grids for the vortex generators and synthetic jets. By reducing the computational cost, the inlet design cycle time can also be significantly reduced. Work on this project has just recently started and we have currently computed the flow for a single micro-vortex generator on a flat plate. Simulations of a synthetic jet on a flat plate and the application of these devices for internal duct flows are in progress.

Future work will include a comparison of the numerical results to experimental data for a better understanding of the flow physics of these devices. Once an understanding of the flow physics is established, the development of computational boundary conditions, which model the flow control devices, can be performed.

This work was done in collaboration with Pieter Buning (NASA Langley).

Efficient and Accurate Solvers for Unsteady Flow Computation

Hester Bijl

The aim is to develop efficient time-dependent approaches for predicting massively separated flows, including automated temporal error control and high-order temporal accuracy. Traditional turbulence models are unable to predict massively separated flows because of the broad spectrum of energy containing eddies present. New methodologies such as Detached Eddy Simulations (DES) have been developed, but require more accurate and efficient time-dependent algorithms.

For this work we will use the TLNS3D code of Veer Vatsa, a Thin Layer Navier-Stokes code with second-order Jameson type spatial discretization. In this code the present strategy for solving time-accurate problems is a second-order backward differencing time integration method with pseudo-time stepping. Solution of the pseudo-time is obtained with an explicit Runge-Kutta method with local time stepping and a fast multigrid solver. We have assessed the temporal and spatial error of this method for a problem for which the exact

solution is known. The temporal error order turned out to be 1.5, exactly between first and second order. This indicates that apart from a zero limit set of points with first-order temporal accuracy, the rest is second order. A similar result was obtained for the spatial error.

Next, we will implement new higher-order unconditionally stable schemes, having precise error control capabilities, in the code. Efficiency and accuracy of these new methods will be assessed, and will be compared with the original time-dependent algorithm. Then we will implement the DES approach into the best of these two time integration methods, i.e., the most efficient and accurate. Finally, the resulting method will be validated for massively separated flows by comparing our numerical results with experiments.

This research was conducted in collaboration with Mark Carpenter (NASA Langley).

Textbook Multigrid Efficiency for the Navier-Stokes Equations

Boris Diskin

A typical modern Reynolds-Averaged Navier-Stokes (RANS) solver (e.g., the CFL3D code) requires approximately 1500 residual evaluations to converge the lift and drag to one percent of their final values for wing-body geometries near transonic cruise conditions. Complex geometry and complex physics simulations generally require many more residual evaluations to converge, and sometimes convergence cannot be attained. On the other hand, it is well-known for fully elliptic problems that solutions can be attained using a full multigrid (FMG) process in far fewer, on the order of 2–4, residual evaluations. This optimal convergence is defined as textbook multigrid efficiency (TME), meaning the solutions to the governing system of equations are attained in a computational work, which is a small (less than ten) multiple of the operation count in the discretized system of equations. Thus, there is a potential gain of more than two orders of magnitude in operation count reduction if TME could be attained for the RANS equation sets. The principal difficulty stems from the fact that the RANS equation sets are a system of coupled nonlinear equations which are not, even for subsonic Mach numbers, fully elliptic, but contain hyperbolic factors. In this period, we systematically studied different versions of the Full Approximation Scheme multigrid algorithm, including distributed relaxation, defect correction, and boundary treatment.

TME has been demonstrated for one of the most basic simulations encountered in fluid dynamics – the incompressible viscous flow past a finite flat plate at high Reynolds number. The flow, although relatively simple, contains several basic elements of the barriers to be overcome in extending textbook efficiencies to the compressible RANS equations, namely entering flows, far wake flows, and boundary layers. The considered equations are the steady, incompressible Navier-Stokes equations in nonconservative form discretized on Cartesian stretched staggered grids. For the numerical calculations, the thin-layer approximation, in which only the viscous terms associated with variations in the coordinate normal to the body are retained, has been used. The foundations for extending TME to generalized coordinates and conservative discretizations have been derived. The approaches to design of efficient relaxation schemes for the compressible flow solvers have been developed and tested for the transonic inviscid flow in a channel. The staggered and colocated arrangements of the discrete unknowns have been investigated.

The next goal ahead is to extend the TME approaches to multidimensional compressible flows.

This research was conducted in collaboration with J.L. Thomas (NASA Langley) and A. Brandt (The Weizmann Institute of Science).

Robust Optimization Including Model Uncertainties and Reliability Constraints

Luc Huyse and R. Michael Lewis

Deterministic optimization improves the performance of a design under given, or assumed, model parameters. However, this is not necessarily an accurate reflection of reality: uncertainties may be associated with the physical model itself (e.g., the parameters in a turbulence model) or the operating conditions may be highly variable (e.g., variable flight speeds and/or angles of attack). Experience has indicated that performing the optimization for discrete design values may result in dramatically inferior design performance when the actual design parameter values are away from those, somewhat arbitrary, design values. The resulting 'optimized' design may prove considerably worse for practical purposes than the original one. The goal is to develop an optimization strategy that explicitly accounts for these uncertainties. The resulting design is robust in the sense that its performance would not be affected dramatically by the uncertainties associated with either the mathematical models themselves or the operating conditions.

A theoretical framework for robust optimization has been developed and its superiority over existing multi-point optimization techniques is demonstrated. Based on a second-order expansion of a mean-value analysis, substantial computational savings should also be achieved when only moderate levels of uncertainty exist for some of the variables or parameters in the mathematical model.

The newly developed techniques are currently applied to an airfoil optimization problem, using the FUN2D code. The minimization of the drag in cruise regime or the maximization of the lift-over-drag ratio is chosen as the objective function. The implementation will be performed in multiple phases: first a robust optimization accounting for the variability in the operating conditions but with deterministic constraints on the design variables will be considered. In a second step, the method will be extended to include the effects of modeling uncertainties on the optimal design as well as reliability-based constraints.

Confidence Bounds on the Hysteresis in Piezoelectric Ceramics

Luc Huyse

An inherent property of piezoelectric materials is the presence of hysteresis and constitutive nonlinearities at moderate to high levels of the driving electric field. To attain the full capabilities of such materials in high performance applications, it is necessary to quantify the variability of these properties. A typical application is the implementation of control design where a prediction of the actual material behavior is required for the full range of operation.

On the basis of experimental results for a PZT-5A compound, a statistical description is developed for the prediction of the randomness associated with the constitutive behavior on the basis of a previously validated domain wall model. Both frequentist and Bayesian techniques are used: the predictor posterior probability distributions for each of the parameters in the hierarchical Bayesian model are compared with traditional confidence interval estimates, based on the Fisher-information matrix and the Maximum Likelihood Estimator.

Future extensions of the model include analysis of the effects of additional physical parameters, such as frequency and temperature, which are currently unaccounted for in the domain wall model. The impact of several production process parameters on the predictor distributions can be included in the analysis as well. The objective is to reduce the variability between different batches and obtain more uniform, or robust, characteristics for PZT actuators.

This work is done in collaboration with Zoubeida Ounaies (ICASE) and Ralph C. Smith (North Carolina State University).

Analysis of Bilevel Approaches to MDO

R. Michael Lewis

Bilevel problem formulations have received considerable attention as an approach to multidisciplinary optimization in engineering. However, careful consideration of the analytical and computational consequences of using bilevel methods has been largely lacking. We have examined the analytical and computational properties of one such approach that has recently received attention under the name collaborative optimization.

Our analysis of collaborative optimization reveals that the resulting system-level optimization problem suffers from inherent computational difficulties due to the bilevel nature of the method, and the stability properties of the resulting system-level optimization problem. Most notably, we have shown that it is impossible to characterize and hence identify solutions of the system-level problems that result in collaborative optimization, because the standard first-order conditions for solutions of constrained optimization problems do not hold. The delinquent analytical features of the system-level problem have serious consequences for the application of conventional nonlinear programming algorithms to finding its solution, and, we believe, explain numerical difficulties with collaborative optimization that have been reported in the MDO literature.

We are currently extending our analysis to a second large class of bilevel approaches. It appears that we can establish fairly general analytical results, which means that many bilevel approaches to MDO unavoidably suffer from serious analytical and computational problems.

This research was conducted in collaboration with Natalia Alexandrov (NASA Langley).

The Coral Project

Josip Lončarić, Thomas W. Crockett, Piyush Mehrotra, Shouben Zhou, and Manuel D. Salas

The cost of developing complex computer components such as CPUs has become so high that scientific applications alone cannot carry the full burden. In the future, scientific computing will have to use mass market leverage to overcome the cost barrier. A cost-effective alternative to high-end supercomputing was pioneered by Beowulf, a cluster of commodity PCs. By now, high performance Beowulf clusters can be built using fast commodity PCs and switched Fast Ethernet. We want to explore the benefits and the limitations of this approach, based on applications of interest to ICASE.

The initial phase of the Coral project, consisting of 32 Pentium II 400 MHz nodes and a dual-CPU server, demonstrated aggregate peak performance in excess of 10 Gflop/s, with sustained performance on CFD applications of about 1.5 Gflop/s. In order to provide a richer environment for further experimentation, a dual-CPU configuration was chosen for the second phase of the Coral project. We have added 32 Pentium III 500 MHz processors (configured as 16 dual-CPU nodes) and two dual-CPU file servers. The resulting system contains 64 compute CPUs with an aggregate of 20 GB of RAM and 440 GB of local disk space. The three dual-CPU servers provide an additional 1.5 GB of RAM and 240 GB of disk space.

We had to resolve a number of performance problems with the new dual-CPU nodes, and tuning continues. Nonetheless, Coral has an excellent price/performance ratio, almost an order of magnitude better than an equivalent proprietary supercomputer design. This conclusion is based on our experience with a variety of applications, ranging from coarse-grained domain decomposition codes to communication-intensive parallel renderers.

For a computer system with this much memory and disk capacity, the ability to move data around within the cluster, as well as to and from other systems, becomes an important consideration. In conjunction with the second phase of the Coral project, we have built a small Gigabit Ethernet testbed which connects the two large file servers, two of the dual-CPU compute nodes, and one of ICASE's graphics workstations.

We were able to reach over 80 percent of the rated bandwidth using 9000 byte jumbo frames, but due to Alteon's hardware/firmware limitations, Gigabit Ethernet latency is about 20 microseconds worse than our Fast Ethernet latency. Gigabit Ethernet based on Alteon's chipset is best suited for bulk data movement. Latency sensitive parallel codes would be better served by another network technology, and we expect to test Giganet products shortly.

The cluster will be used to develop and run research codes of interest to ICASE and NASA Langley, and to evaluate price/performance tradeoffs among various hardware, software, and networking configurations.

The Coral Project web site: `http://www.icase.edu/CoralProject.html`.

Active Shielding and Control of Environmental Noise

Josip Lončarić, Victor S. Ryaben'kii, and Semyon V. Tsynkov

Rejection of exterior noise caused by periodic sources such as propellers or turbines would significantly enhance passenger comfort and reduce noise fatigue on long flights. Passive sound absorbing materials help at high frequencies, but to be effective below about 1 kHz their weight penalty becomes significant. Active noise control can reduce low frequency noise with less weight penalty. We present the mathematical foundations of a new active technique for control of the time-harmonic acoustic disturbances.

Unlike many existing methodologies, the new approach provides for the exact volumetric cancellation of the unwanted noise in a given predetermined region of space while leaving those components of the total sound field deemed as friendly unaltered in the same region. Besides, the analysis allows us to conclude that to eliminate the unwanted component of the acoustic field in a given area, one needs to know relatively little; in particular, neither the locations nor the structure of the external noise sources need to be known. We constructed the general solution for the aforementioned noise control problem. The apparatus used for deriving the general solution is closely connected to the concepts of generalized potentials and boundary projections of Calderon's type. To prove that the new technique is appropriate, we thoroughly worked out a simple two-dimensional model example that allows full analytical consideration, including optimalization of the control effort.

In order to develop numerically computable solutions, we plan to describe the discrete framework for the noise control problem parallel to the continuous one. This discrete framework is obtained using a difference potentials method; in the future it is going to be used for analyzing complex configurations that originate from practical designs. Once we have computed the solution for a particular configuration, we intend to investigate the possibilities of optimizing it according to the different criteria that would fit different practical requirements. We expect to discuss the applicability of the technique to quasi-stationary problems, future extensions to the cases of the broad-band spectra of disturbances, as well as other possible applications which may include different physics, such as electrodynamics, and different formulations of the boundary-value problems, such as scattering.

Large Eddy Simulation Using a Parallel Multigrid Solver

Dimitri J. Mavriplis

The failure to develop a universally valid turbulence model coupled with recent advances in computational technology have generated a greater interest in the large-eddy simulation approach for computing flows with large amounts of separation. This approach involves resolving the large-scale unsteady turbulent eddies down to a universally valid range in the hope of yielding a more generally valid simulation tool. The purpose

of this work is to develop a large-eddy simulation capability based on an existing unstructured grid Navier-Stokes solver. The use of unstructured grids, which facilitates the discretization of complex geometries and adaptive meshing techniques, is expected to enhance the flexibility of the resulting simulation capability.

An unsteady Reynolds-averaged Navier-Stokes flow solver based on unstructured meshes has been developed and is currently undergoing validation. The initial validation case consists of the unsteady flow over a circular cylinder, with emphasis on obtaining the correct Strouhal number. A grid convergence study (in time and space) has been performed in order to determine the resolution requirements for a given level of accuracy. The second test case, currently underway involves the post-stall flow over a three-dimensional wing geometry. These computations are being performed on the ICASE PC cluster, Coral.

Once a validated unsteady Reynolds-averaged Navier-Stokes solver is in hand, the implementation of a subgrid scale model will be pursued. This model will be validated by computing several standard decaying turbulence test cases, and will then be applied to cases of aerodynamic interest.

This work is being carried out in collaboration with Juan Pelaez (Old Dominion University and ICASE).

Scalability of an Unstructured Multigrid Flow Solver Using Various Programming Models

Dimitri J. Mavriplis

Under the ASCI program, the need to simulate very large problems on massively parallel machines has focused attention on the relative merits of various parallelization strategies. The implementations of interest include the message-passing interface (MPI) library and the OpenMP parallelization standard. While OpenMP has been devised for use on shared-memory architectures, MPI can be implemented on any type of architecture, but is primarily attractive for distributed memory architectures. Furthermore, the emergence of clusters of shared memory multi-processors as a viable scientific computing architecture has resulted in interest in hybrid OpenMP-MPI programming models.

The goal of this work is to develop a parallel unstructured mesh flow solver using both MPI and OpenMP, and to examine the scalability of the solver on various architectures using the different programming models either alone, or in the combined hybrid mode.

An existing unstructured multigrid flow solver, which was originally parallelized using the MPI approach, has been upgraded to include an OpenMP parallelization strategy. OpenMP parallelization is achieved within the existing MPI parallel structure, in a two-level nested fashion. In this manner, the same code can be run in pure MPI mode, in pure OpenMP mode, or in the combined OpenMP-MPI mode. Two approaches for overlapping shared memory and distributed memory communication have also been implemented.

Extensive benchmarks have shown MPI and OpenMP to yield equivalent results on a variety of architectures including the CRAY SV1, the SGI Origin 2000, and dual-CPU Pentium PCs, when either programming model is used exclusively. On the other hand, combined OpenMP/MPI benchmarks on the SGI Origin 2000 have shown performance degradation when a large number of OpenMP threads are specified under MPI.

While these results indicate that hybrid parallel programming models suffer on the Origin 2000, such models may be more suitable for clusters of shared memory multi-processor nodes, or relatively powerful multi-processors linked together over the internet. Future work will involve benchmarking the current implementation on new architectures as they become available.

Algorithms for the Solution of Radiation Diffusion Problems

Dimitri J. Mavriplis

Under the ASCI program, the simulation of radiation transport phenomena has been identified as one of the most time consuming elements within large simulation codes. The objective of this research is to

investigate the effectiveness of unstructured multigrid algorithms in efficiently solving unsteady radiation transport problems, in the radiation diffusion limit.

A two-dimensional unstructured grid radiation diffusion solver has been developed, which employs agglomeration multigrid to converge the nonlinear problem at each physical time step. In a first approach, a nonlinear FAS multigrid scheme is used to directly converge the nonlinear problem at each time step, whereas in the second approach, a Newton iteration scheme is applied to the nonlinear problem and a linear multigrid algorithm is employed to solve the linear system at each Newton iteration.

A complete validation of the solver in terms of accuracy is to be completed in the near future. Thereafter, a detailed study of the advantages of linear multigrid versus nonlinear multigrid will be investigated for a range of problems.

High-order Finite Difference Approximations for Linear Problems

Jan Nördstrom

The objective of this research is to develop difference operators that guarantee strict stability for PDE's with varying coefficients on curvilinear meshes.

The effects of reformulating the difference operation, varying the norm and adding artificial viscosity, were considered.

We plan to continue the analysis mentioned above.

This research was conducted in collaboration with Mark H. Carpenter (NASA Langley).

Global Artificial Boundary Conditions for Aerodynamic and Aeroacoustic Computations

Viktor Ryaben'kii and Semyon Tsynkov

Many problems in scientific computing, including those that present immediate practical interest, e.g., problems of acoustic radiation/propagation/scattering, are formulated on infinite domains. Therefore, any numerical methodology for solving such problems has to be supplemented (or, rather, preceded) by some technique that would lead to a finite discretization. Typically, the original domain is truncated prior to the actual discretization and numerical solution. Subsequently, one can construct a finite discretization on the new bounded computational domain using one of the standard techniques: finite differences, finite elements, or other. However, both the continuous problem on the truncated domain and its discrete counterpart will be subdefinite unless supplemented by the appropriate closing procedure at the external computational boundary. This is done by using artificial boundary conditions (ABCs); the word "artificial" emphasizing here that these boundary conditions are necessitated by numerics and do not come from the original physical formulation.

At the current stage of the aforementioned project, we are focusing on the development of highly accurate unsteady ABCs for the numerical simulation of waves propagating with finite speed over infinite domains. The major difficulty is the nonlocal character of these boundary conditions in both space and time. Typically, the spatial nonlocality is tolerable, while the temporal one is prohibitively expensive. To overcome this difficulty, we make use of the presence of lacunae (in the sense of Petrovsky) in the solutions of some linear hyperbolic PDEs in oddly dimensional spaces. Based on this property of the solutions, we construct a special lacunae-based algorithm for the long-term numerical integration of hyperbolic PDEs. The algorithm can employ any convergent finite-difference scheme as a core computational technique; then the standard computational procedure of the scheme is modified so that to guarantee a temporally uniform grid convergence for arbitrarily long time intervals. This is a unique distinction of the new algorithm unparalleled by other

methods. The algorithm obviously does not accumulate error in the course of time and besides, it has fixed non-growing expenses per time step. Incorporating this lacunae-based algorithm into the general procedure of constructing the ABCs based on Calderon's projections and difference potentials, we gain the restriction of the temporal nonlocality of the boundary conditions. Some initial calculations for the three-dimensional wave equation excited by a compactly supported stationary source have already provided the encouraging results. The code for the wave equation excited by a moving source is in the development stage. The paper "Long-time Numerical Computation of Wave-type Solutions Driven by Moving Sources" by V. Ryaben'kii, S. Tsynkov, and V. Turchaninov has been submitted for publication to Applied Numerical Mathematics.

Future research in the framework of this project will primarily concentrate on the development of unsteady ABCs for the actual problems in acoustics, including the advective case, and electromagnetics.

The project is a collaborative effort with V. Turchaninov (Keldysh Institute for Applied Mathematics, Russian Academy of Sciences). The project is supported by the Director's Discretionary Fund.

High-order Discontinuous Galerkin Method and WENO Schemes

Chi-Wang Shu

Our motivation is to have high-order non-oscillatory methods for structured and unstructured meshes, which are easy to implement for parallel machines. The objective is to develop and apply high-order discontinuous Galerkin finite element methods and weighted ENO schemes for convection dominated problems. The applications will be problems in aeroacoustics and other time-dependent problems with complicated solution structure.

Jointly with Harold Atkins at NASA Langley, we are continuing in the investigation of developing the discontinuous Galerkin method to solve the convection-dominated convection diffusion equations. We have achieved preliminary results on a potentially very effective local preconditioner which may help in designing efficient implicit time discretization techniques for discontinuous Galerkin methods for convection-diffusion problems. Jointly with Sigal Gottlieb and Eitan Tadmor, we have reviewed and further developed a class of strong stability preserving (SSP) high-order time discretizations for semi-discrete method of lines approximations of partial differential equations. Termed TVD (total variation diminishing) time discretizations before, this class of high-order time discretization methods preserves the strong stability properties of first-order Euler time stepping and has proved very useful especially in solving hyperbolic partial differential equations. The new developments include the optimal explicit SSP linear Runge-Kutta methods, their application to the strong stability of coercive approximations, a systematic study of explicit SSP multi-step methods, and the study of the strong stability preserving property of implicit Runge-Kutta and multi-step methods.

Research will be continued for high-order discontinuous Galerkin methods and weighted ENO methods and their applications.

On the Application of Global Boundary Conditions with the New Generation of Flow Solvers

David Sidilkover

Previous work indicates that the global highly accurate artificial boundary conditions implemented along with a multigrid flow solver may lead to major improvements in the overall algorithmic efficiency. These improvements are due to the substantial reduction of the domain size with no accuracy loss, better robustness, and faster multigrid convergence. However, the flow solvers considered so far were the standard ones which are far from being optimally efficient. The objective of this research is to study the application of the global boundary conditions with the new generation of the flow solvers, which are optimal.

The new generation of flow solvers is based on factorizing the flow equations, i.e., distinguishing between the factors of different type. This simplifies the derivation of the global boundary conditions, since each factor can be treated separately for this purpose. The boundary condition accurately takes into account the structure of the far-field solution; for simple settings this is done directly using Fourier transform and separation of variables, for more general settings the difference potentials method (DPM) is used. The issue of relaxing properly the discrete equations together with the new global boundary conditions requires special attention in order to maintain the optimal convergence rates of the overall solver. This was the main subject of the current research. As an outcome of the project a multifold reduction of the size of the computational domain has been demonstrated, while preserving the accuracy of the solution and the optimal convergence rates of the solver.

Currently, the algorithm based on the separation of the elliptic equation for the pressure is under study. In the future we plan on applying the global boundary conditions in conjunction with the new upwind high-resolution factorizable (UHF) scheme.

This is joint work with S.V. Tsynkov (Tel Aviv University and ICASE) and T.W. Roberts (NASA Langley).

Upwind High-resolution Factorizable (UHF) Schemes for the Equations of Fluid Flow
David Sidilkover

In order to develop an optimally efficient solver for the general equations of fluid flow a new discretization scheme is needed. It has to imitate the mixed type of the PDEs, i.e., to be *factorizable* – to make the different factors of the equations distinguishable on the discrete level. In addition to that it has to possess the shock-capturing property, i.e., to be conservative and high-resolution. The scheme also has to be h-elliptic, i.e., its residuals should be sensitive to the high-frequency error components. There was no discretization that satisfies all these requirements. The main objective of this research was to develop one.

An upwind high-resolution factorizable (UHF) scheme for the Euler equations in two dimensions has been constructed. The factorizability property can be utilized for the purpose of designing an optimal solver through a "projection-type" relaxation. This relaxation procedure relies on the auxiliary potential and stream-function variables. Another important implication of the factorizability property is that the scheme should not lose accuracy for the low Mach number flow. In other words, the proposed approach essentially provides the long-awaited unification of the compressible and incompressible flow solvers. The subsonic version of the scheme was previously constructed. The current work is concerned with extending these ideas to the transonic/supersonic regimes. The main issue – to acquire the shock-capturing capability while preserving the *factorizability* property – has been resolved.

Some of the current work is devoted to extending the scheme/solver to the three-dimensional case and general body-fitted grids.

PHYSICAL SCIENCES

Parallelization of the Jet Simulation Code

Ayodeji Demuren

The computer code used for direct numerical and large eddy simulation of complex 3D jets is to be parallelized to run on a multitude of processors. Scalability, portability and accuracy are considered of paramount importance. The goal is to be able to run the code on any available multi-processor system, such as workstation clusters, Beowulf systems, or SMP systems from CRAY, SGI, HP, or SUN.

The chosen approach was to convert a previously developed higher-order-accurate numerical formulation for the simulation of complex jets to enable parallel simulation via domain decomposition. MPI library calls are used for communication. Stability of the multi-domain simulation is assured by use of difference operators, which satisfy summation-by-parts (SBP) criteria. Parallelization is completed for explicit finite-difference SBP schemes. Performance tests are being carried out. In addition, the original compact schemes are also to be parallelized because of their advantage in eliminating odd-even decoupling problems in incompressible flow simulation.

We plan to complete parallelization tasks and then compare performance of compact and explicit SBP codes.

Laser-induced Thermal Acoustics as a Practical Optical Diagnostic

Roger Hart

Laser-Induced Thermal Acoustics (LITA) is a relatively new optical diagnostic technique that shows great promise to extend the capabilities of non-intrusive measurement beyond those of well-established techniques such as Laser Doppler Velocimetry (LDV) and Particle Interval Velocimetry (PIV). Accurate, practical, non-intrusive measurement of flow parameters such as velocity, temperature, and pressure is of considerable interest to the aerodynamics community. Possible applications for such measurements range from freestream calibration and turbulence measurement to studies of flow separation and measurement of higher-order structure functions for the validation of small-scale turbulence models. One particular advantage of LITA over LDV and PIV is that it is seedless. One disadvantage relative to PIV is that extension of LITA measurements to a plane may not be feasible (although measurements along a line may prove practical).

Two lasers are employed in LITA. A short-pulse pump laser creates a pair of acoustic wavepackets in the air flow. Light from a probe laser is Bragg-scattered from the wavepackets to a detector; the scattered light is Doppler shifted due to the motion of the wave packets. Analysis of the spectral content of the demodulated signal allows determination of local speed of sound (and, thus, temperature), flow velocity, and pressure (from the decay rate of the sound waves). Effort during the current reporting period has concentrated entirely on the design and construction of a fieldable prototype LITA apparatus, which will be evaluated at the Langley Basic Aerodynamic Research Tunnel (BART) in a test beginning mid-August 2000. The assembly of the transmitter portion of the apparatus is essentially complete as of this writing. Work continues on the receiver. Previous laboratory demonstrations of LITA thermometry and velocimetry have employed laboratory-style lasers and hardware which emphasize flexibility at the expense of compactness, stability, and 'user-friendliness'. The apparatus currently under construction is far smaller and should require far less frequent adjustment to optics and lasers while providing performance equal to or exceeding that previously demonstrated.

Near-term activities include: completion of the new apparatus (anticipated by mid-April); extensive laboratory testing and evaluation of the new hardware; and development and testing of new data analysis software. All these should be largely complete before the beginning of the BART test. Possible follow-on activities are currently under discussion. One long-term goal is to develop an apparatus stable enough and reliable enough to operate in the plenum at NTF. The current prototype is intended to be a step in that direction.

This research is conducted in collaboration with Jeff Balla and Greg Herring (NASA Langley).

Confidence Bounds on the Hysteresis in Piezoelectric Ceramics

Luc Huyse

An inherent property of piezoelectric materials is the presence of hysteresis and constitutive nonlinearities at moderate to high levels of the driving electric field. To attain the full capabilities of such materials in high performance applications, it is necessary to quantify the variability of these properties. A typical application is the implementation of control design where a prediction of the actual material behavior is required for the full range of operation.

On the basis of experimental results for a PZT-5A compound, a statistical description is developed for the prediction of the randomness associated with the constitutive behavior on the basis of a previously validated domain wall model. Both frequentist and Bayesian techniques are used: the predictor posterior probability distributions for each of the parameters in the hierarchical Bayesian model are compared with traditional confidence interval estimates, based on the Fisher-information matrix and the Maximum Likelihood Estimator.

Future extensions of the model include analysis of the effects of additional physical parameters, such as frequency and temperature, which are currently unaccounted for in the domain wall model. The impact of several production process parameters on the predictor distributions can be included in the analysis as well. The objective is to reduce the variability between different batches and obtain more uniform, or robust, characteristics for PZT actuators.

This work is done in collaboration with Zoubeida Ounaies (ICASE) and Ralph C. Smith (North Carolina State University).

Analysis of Lattice Boltzmann Equation: Dispersion, Dissipation, Isotropy, Galilean Invariance, and Stability

Li-Shi Luo

The method of the lattice Boltzmann equation (LBE) has been applied to various areas in computational fluid dynamics. The dynamics of the lattice Boltzmann equation evolves on a highly symmetric lattice space, which is usually square in 2D and cubic in 3D. Therefore, the transport properties of the LBE method are dictated by the symmetry of the underlying lattice. The present work proposes a general procedure to study the dispersion, dissipation, isotropy, Galilean invariance, and stability of the LBE method.

We first propose a generalized 9-velocity LBE model in two-dimensional space. The model is constructed in moment space instead of velocity space, and possesses a maximum number of adjustable parameters allowed by the freedom of the LBE model. We then analyze the dispersion equation of the linearized evolution operator to obtain the generalized hydrodynamics of the model (the wave vector dependence of the transport coefficients). We study the dispersion, dissipation, isotropy, Galilean invariance, and stability of the LBE model, and optimize the properties of the model by tuning the adjustable parameters in the

model according to our linear analysis. We find that the proposed model is superior to the popular BGK LBE model in terms of stability, isotropy, and Galilean invariance. Various LBE models are also analyzed and compared with each other.

We intend to extend our work to 3D LBE models or other more complicated LBE models, such as LBE models for visco-elastic fluids.

The present work is a result of collaboration with Prof. Pierre Lallemand (Laboratoire ASCI-CNRS, Université Paris-Sud (Paris XI Orsay), France). The present work has been partly funded by NASA Langley Research Center under the program of "Innovative Algorithms for Aerospace Engineering Analysis and Optimization."

Networked Rectenna Array for Smart Material Actuators

Yeon-Gon Mo

The concept of microwave-driven smart material actuators is envisioned as the best option to alleviate the complexity associated with hard-wired control circuitry. Networked rectenna array receives microwave power and converts it into DC power for an array of smart actuators. To use microwave power effectively, the concepts of power allocation and distribution (PAD) circuit was adopted for networking rectenna array. The PAD circuit is imbedded into a single embodiment of rectenna and actuator arrays. The thin-film microcircuit embodiment of PAD circuit adds an insignificant amount of rigidity to membrane flexibility.

A 3x3 bow-tie rectenna array was characterized for various input conditions, such as transmitted power, frequency, beam incident angle, and beam polarization. DC power generated by the bow-tie rectenna array was collected under different conditions. In order to demonstrate a concept of PAD, relay switching circuit, IR remote control circuit, and a 3x3 bow-tie rectenna array were integrated. The control of power distribution was conducted by a relay switching circuit that included 15 relays. Also, a bypassing unit was incorporated into each DC power source to protect against a power failure. All switching was controlled by an IR remote control unit that adopted 38 kHz as a transmitter carrier wave frequency. The networked actuators with 3x3 bow-tie rectenna arrays were tested to correlate the network coupling effect, power allocation and distribution, and response time. Also, for better control of actuators, a circuit design was done. The features of the design are 16-channel computer control of actuators by a PCI board and the compensator for a power failure of one or more rectennas. The test work performed thus far shows that the approach works well.

Further investigation will be extended to design and test a PAD circuit system with higher nodal elements.

High Temperature Piezoelectric Polyimides

Zoubeida Ounaies

High performance piezoelectric polymers are of interest to NASA as they may be useful for a variety of aerospace and aircraft applications. Amorphous polyimides containing polar functional groups have been synthesized and investigated for potential use as high temperature piezoelectric sensors. The intent of this project is to elucidate the mechanism and key components required for developing piezoelectricity in amorphous polymers and further to apply this understanding in designing a novel high temperature piezoelectric polyimide.

The effect of structural changes, including variations in the nature and concentration of dipolar groups, on the piezoelectric behavior is examined. The remanent polarization, the dielectric relaxation strength and the various piezoelectric coefficients are reported. The thermal stability of the piezoelectric effect of one

polyimide was evaluated as a function of various curing and poling conditions under dynamic and static thermal stimuli. The piezoelectricity of this polymer was stable under both dynamic and static thermal stimuli and statistically no loss was observed up to 100°C. Since this amorphous piezoelectric polyimide can generate a piezoelectric response at elevated temperatures, its thermal stability should be beneficial for high temperature aerospace applications.

Fundamental structure-piezoelectric property insight will enable the molecular design of polymers possessing distinct improvements over state-of-the-art piezoelectric polymers including enhanced polarization, polarization stability at elevated temperatures, and improved processability.

This work is done in collaboration with Joycelyn S. Harrison (NASA Langley) and Cheol Park (National Research Council).

Piezoelectric Ceramics for Use as Actuators

Zoubeida Ounaies

The intent of this project is to assess the electrical characteristics of piezoelectric wafers for use in aeronautical applications such as active noise control in aircraft. The determination of the capacitive behavior and power consumption of these actuators is necessary to optimize the system configuration and to design efficient driving electronics. Power consumed by the lead zirconate titanate (PZT) is the rate of energy required to excite the piezoelectric system along with power dissipated due to dielectric loss and mechanical and structural damping. Overall power consumption, thus, needs to be quantified as a function of peak applied voltage and driving frequency.

The capacitance and loss tangent of a piezoelectric wafer were characterized as a function of driving field and frequency. An analytical model based on RC series circuit was used to estimate the capacitance and resistance of the piezoelectric wafer. Mathematical relations were then developed to relate the capacitance and resistance values to the peak driving voltage and frequency. It was demonstrated that by incorporating the variation of capacitance and power loss with voltage and frequency, satisfactory estimates of power requirements can be obtained. These relations allow general guidelines in selection and application of piezoelectric actuators and driving electronics for active control applications.

In future studies, the electromechanical efficiency of the piezoelectric actuators will be investigated in this frequency range.

This work is done in collaboration with Thomas L. Jordan (NASA Langley).

A Hysteresis Model for Piezoelectric Ceramics

Zoubeida Ounaies

This project involves characterizing the hysteresis behavior of piezoelectric ceramics, both computationally and experimentally, under moderate to high fields in order to estimate their nonlinear behavior. This can be useful in the implementation of control design where the prediction of the behavior of the actuator throughout the full range of operation is necessary. While quasistatic models are suitable for initial material characterization in some applications, the reduction in coercive field and polarization values which occurs as frequencies increase must be accommodated to attain the full capabilities of the materials.

Towards that goal, hysteresis data was taken at frequencies ranging from 0.1 Hz to 5 Hz to quantify the frequency effects of the switching mechanism. The model was modified to address the reduction in polarization which occurs as frequencies increase from quasistatic levels to low frequency regimes. This is modeled by determining the probability that dipoles achieve the energy required to overcome energy barriers

and switch orientation when an external field is applied. The model was found to accurately quantify the data at 0.1 Hz, which is the regime in which the parameters were estimated. The model also accurately predicts the decrease in polarization at 1 Hz, but does not yet incorporate the decrease in the coercive field, which occurs as frequency increases.

In its current formulation, the model is restricted to low frequency regimes since it does not yet incorporate rate-dependent mechanisms for quantifying the energy required to reorient dipoles when translating domain walls. The extension of the model to include these mechanisms is necessary in order to increase the applicable frequency range of the model and is under current investigation.

This work is lead by Ralph C. Smith (North Carolina State University).

Molecular Dynamics Simulations for Mechanical and Electrical Properties of Metallic and Polymer Systems

Sun M. Paik

For a quantitative understanding of the underlying mechanisms for mechanical and electrical failures it is necessary to investigate and quantify the effects of atomistic structures and defects on material properties. The Molecular Dynamics (MD) simulation technique is an exact (within numerical accuracy of the computer) deterministic numerical simulation method, and one of the best-suited methods to study local structures of a system. In this investigation, we will study the mechanical properties of metallic systems and the dielectric properties of polymers using the Molecular Dynamics simulation technique providing failure sensing and preventing mechanisms.

The materials simulations fall into three different categories: (1) deterministic simulations such as the MD simulation, (2) stochastic such as a Monte Carlo (MC) simulation, and (3) *ab initio* calculation. These methods have both *pros* and *cons*, but the MD method is best suited for our purpose. Using the MD technique we will simulate a stable structure of metallic alloy systems and effects of the structure factors like alloy composition ratio, segregation, grain boundaries, impurities, and defects on crack formation and propagation. A many-body Embedded Atomic Method (EAM) potential is used for the interatomic interactions between the metal atoms. We have completed coding the MD program in C for a single metallic (Au) system, and are extending to alloy systems. For polymer systems, we are in the stage of developing a suitable empirical molecular interaction potential.

In the future we will simulate and calculate various mechanical and dielectric properties of more complicated metallic alloy and polymer systems. We also plan to write MD code for a parallel computer.

This work is done in collaboration with Dr. Min Namkung (NASA Langley).

Propagation of Acoustic Pulses in Non-uniform Flows

Alex Povitsky

Many practical problems in aeroacoustics that involve noise generation and propagation are actually problems of propagation of disturbances in a non-uniform mean flow. Vorticity, entropy, and sound disturbances are typical building blocks of the turbulence in an incompressible flow, of temperature fluctuations in low-speed turbulent heat transfer, and of conventional sound, respectively. Available previous studies either consider propagation of pulses in uniform flows or split the disturbance velocity on three aforementioned components to solve the corresponding equations separately.

In this study, we simulate numerical propagation of acoustic, vortical, and entropy dipole and monopole pulses in a non-uniform incompressible stagnation mean flow. Performing higher-order numerical simulation

of Euler equations for disturbance variables, we found that baroclinic generation of disturbance vorticity appears to be a key process in energy transfer between a non-uniform mean flow and a propagating disturbance. These phenomena lead to amplification of sound waves originated from aforementioned pulses. In turn, Lighthill's equation for the total enthalpy appears to be mutually dependent with the vorticity equation that should be accounted in existing computational aeroacousticsmethodologies.

In our future research we shall study the mechanism of amplification of disturbances in non-uniform flows near the leading edge of a wing. Another direction of investigation includes generation of sound and its propagation in a non-uniform flow while a vortex travels through a system of oblique shock waves and expansion waves.

Optimization of Mixing Jet Reactor for Production of Carbon Nanotubes
Alex Povitsky

Single-walled carbon nanotubes are produced from CO with the help of a catalyst. Particles of the catalyst get into the reaction chamber with the central CO jet. The particles must be heated from 200°C (central jet temperature) to temperatures close to 1000°C (reactor temperature) as rapidly as possible. Slowly heated catalytic particles quickly lose their catalytic ability. The mixing reactor studied here consists of an arrangement of jets designed to heat up catalytic particles carried by a turbulent jet issuing from a round nozzle. This jet entrains the hot gas within the reactor cavity and heats up. However, this heating is considered too slow to produce a significant amount of carbon nanotubes. To increase mixing, hot turbulent jets issuing from several nozzles surrounding the central nozzle are used. These jets are inclined to the central jet to further increase the mixing rate.

A 3D code has been developed in the course of this study to solve a turbulent recirculating flow with variable density. The numerical algorithm is based on finite-volume discretization and Patankar's SIMPLER methodology to split momentum and continuity equations. Generation of turbulence in co-flowing jets, generation of vortices and rotation of highly inclined jets, and optimization of the number of jets with the constraint on the total mass flux were studied systematically. Numerical simulations show that a pair of peripheral jets inclined to the central jet produces the maximum heat-up of catalytic particles. The fast mixing and heating of the central jet occurs in cases with strong angular non-uniformity of the merged jet and strong spreading of particles. Numerical results are used to design an experimental reactor to produce carbon nanotubes for Professor R. Smalley of Rice University.

Our future research includes the study of counter-flowing jets, pulsed jets, and initial rotation of peripheral jets. We also plan to incorporate compressibility and more general reactor geometry in our computer code. Then we shall study compressible speed-up of jets for fast heating, aerodynamic heating by interaction of counter-flowing jets, and optimization of conical (shower-head) geometry of the reactor.

Stochastic Description of Transitional Flows
Robert Rubinstein

There is still no well-established procedure for incorporating transition in turbulence calculations. While aerodynamic flows can be computed successfully using any of several different turbulence models if the transition location is prescribed in advance, no single turbulence model can reliably predict transition. If transition is computed incorrectly, the entire flow calculation is generally unsatisfactory.

The previous work on stochastic receptivity has been continued into the phase of weakly nonlinear growth of disturbances by deriving the evolution equation for the joint probability density function of resonant triads

of Tollmien-Schlichting waves. Numerical integration of this equation shows that phase correlations among the modes develop, even if the phases are initially uniformly distributed. By computing some representative cases, it has been verified that the stochastic receptivity equations predict that the initial disturbances will, in fact, have random phase. Mode amplitudes show the opposite trend, namely a tendency to become distributed over a large range of values even if the initial amplitude is nearly deterministic.

Work has begun on a stochastic formulation of the parabolized stability equations. These equations are believed to offer a more general and realistic description of transition than the theory of resonant wave interactions.

This work was performed in collaboration with Meelan Choudhari (NASA Langley).

Electroactive Polymers for Lightweight, High Performance Actuators

Ji Su

Electromechanical coupling effects such as piezoelectricity and electrostriction are widely used in many transducer and actuator technologies. There are three material classes of electromechanically active materials: ceramics, polymers, and ceramic-polymer composites. Polymers have the distinct advantage of being lightweight, conformable, tough, and easily processed. Poly(vinylidene fluoride) (PVDF) is currently the most widely used and only commercially available electroactive polymer. This material has excellent sensing capability but is quite limited as an actuator. Within the last decade, research to develop polymers with high actuation capability has resulted in the development of electrostrictive polyurethane elastomers. Electromechanical strain in polyurethane has been observed as high as 4%. However, the polyurethane elastomer has low mechanical modulus and consequently, very little actuation capability. The object of this project is to develop new electrostrictive polymers that can offer not only large electric field induced strain, but also mechanical force or mechanical energy density, therefore desired actuation authority.

Recently, researchers in the Advanced Materials and Processing Branch (AMPB) at LaRC have discovered electrostrictive response in a new class of graft-elastomer polymers. Electrostrictive strain in excess of 4% has been observed in the elastomer. The mechanical force output and specific mechanical energy density of the electrostrictive graft elastomer can be more than one order of magnitude higher than that of electrostrictive polyurethane. However, the large electric field induced strain requires tougher and more flexible electrode materials since the conventional metal electrode materials do not maintain functionality for a long time under large cycle strain. A conductive polymer, polypyrrole, was selected as a candidate to replace metal electrodes for the electrostrictive graft elastomer. A one-step polymerization-deposition method was employed in developing a conductive polymer-coated electrostrictive graft elastomer sandwich system. In this procedure, anthraquinone-2-sulfonic acid, sodium salt monohydrate, 5-sulfosalicylic acid dihydrate, and ferric chloride hexahydrate were dissolved in distilled water for solution one. An aqueous solution of pyrrole was poured into solution one. Electrostrictive graft elastomer films fixed on a metal frame were then immediately immersed into the polymerizing pyrrole solution. The opposing surfaces of the graft elastomer films were coated with polypyrrole. The conductive polymer electroded electrostrictive graft elastomer films were produced. The resultant polypyrrole electrodes were smooth, coherent, and adhered well to the graft elastomer and exhibited acceptable electric conductivity.

Characterization of the mechanical, dielectric, and electromechanical properties of the conductive polymer-coated electrostrictive graft elastomer and fabrication of a prototype bending actuator using the all polymer electrostrictive systems are the planned work. Property comparison between the metal-coated graft elastomer films and the conductive polymer-coated graft elastomer films will also be conducted.

This research was conducted in collaboration with Joycelyn S. Harrison (NASA Langley).

COMPUTER SCIENCE

Parallel Implicit Solvers for Simulation of Multiscale Phenomena

David E. Keyes

The development and application of parallel implicit solvers for multiscale phenomena governed by PDEs are our chief objectives. Newton-Krylov-Schwarz (NKS) methods have proven to be broadly applicable, architecturally versatile, and tunable for high performance on today's high-end commercial parallel platforms (e.g., Cray T3E, SGI Origin, IBM SP). Both structured-grid and unstructured-grid CFD legacy codes have been ported to such platforms and reasonable objectives for algorithmic convergence rate, parallel efficiency, and raw floating point performance have been met. However, architectural challenges have increased on the next generation of high-end machines, as represented, for instance, by the ASCI "blue" machines at Lawrence Livermore and Los Alamos National Laboratories, and also on Beowulf clusters, such as ICASE's "Coral." Our primary efforts are concentrated on algorithmic adaptations of NKS methodology appropriate for the emerging architectures and on evaluation of new software tools and methodology to get the most performance out of them.

The general approach embodied in the NKS family of algorithms is documented in previous ICASE technical reports, among other places. Newton's method, robustified by pseudo-transient continuation, generates global linear systems that are solved in a matrix-free manner by Krylov iteration, and preconditioned locally. In the absence of timestepping, a non-local two-level preconditioner based on a coarsening of the problem is also required to maintain algorithmic scalability at a modest cost to asymptotic parallel efficiency, but most nonlinear problems require some form of timestepping, even for the computations of steady states.

In the most recent reporting period, our work was recognized with one of the 1999 Bell Prizes (awarded in November 1999 at Supercomputing). Our implementation of the NASA code FUN3D ran in a scalable way on a fixed-size problem of approximately 11 million degrees of freedom on up to 3,072 dual-processor nodes of the ASCI Red machine (made up of Intel Pentium Pro processors), achieving a total of 227 Gigaflop/s. Multithreading was applied to the flux computations, but not yet to the linear algebra subroutines. The flux computations amount to about half of the total runtime before multithreading, and approximately 2,000 "work units" of global flux evaluations are required to achieve ten orders of magnitude reduction in the initial residual.

We will continue to develop NKS methods in implicit parallel CFD and radiation transport, examining a variety of algorithmic, programming paradigm, and architectural issues. We will also use the parallel PDE analysis capability as part of a PDE-constrained optimization framework, known as Lagrange-Newton-Krylov-Schwarz.

This work is joint with W. Kyle Anderson (NASA Langley), Dana Knoll (Los Alamos National Laboratory), Dinesh Kaushik (Old Dominion University and ICASE), and Satish Balay, William D. Gropp, Lois C. McInnes, and Barry F. Smith (Argonne National Laboratory).

Model Checking is Refinement

Gerald Lüttgen

This work is inspired by NASA's search for a suitable *design framework* for mode logics of flight guidance systems. This framework should support the transitioning from the requirements phase to the design phase and enable the efficient analysis of mode logics. We have previously shown that *model checking* is an

efficient technique for reasoning about mode logics, which are operationally specified by state machines and where system properties are stated by temporal logic formulas. The objective of the present research is to develop a semantic foundation for *heterogeneous specification languages* combining labeled transition systems and linear-time temporal logic (LTL). The setting is required to support a compositional notion of stepwise refinement, which allows for the gradual replacement of temporal logic formulas by labeled transition systems satisfying the formulas.

The considered semantic domain are *Büchi automata*, for which we developed and studied two *refinement preorders*, called Büchi may-preorder and Büchi must-preorder, respectively. These preorders, for which also alternative characterizations were established, conservatively extend preorders introduced in DeNicola and Hennessy's testing theory from labeled transition systems to Büchi automata. We showed that LTL model checking coincides with preorder checking for the Büchi must-preorder, when LTL formulas are translated into purely nondeterministic Büchi automata. Thereby, a refinement relation was obtained, which is compatible with the LTL satisfaction relation while extending a traditional refinement preorder that has proved to be compositional and fully abstract in most process–algebraic settings. In particular, our results guarantee that system properties are preserved when gradually transitioning from system specifications to system designs.

Future work will focus on developing *concrete languages* mixing assertional and operational styles of specification, whose semantics can be given in terms of Büchi automata. For these languages, we will then study compositionality issues, fully abstractness, and axiomatic characterizations of our refinement preorder. We will also investigate how our technique can be integrated in *design tools*, which are expected to be used for the development of future flight guidance systems.

This work is done in collaboration with Rance Cleaveland (SUNY at Stony Brook).

Fully Abstract Statecharts Semantics via Intuitionistic Kripke Models
Gerald Lüttgen

Statecharts is a visual notation for specifying reactive systems, which extends state machines by concepts of concurrency and hierarchy. Unfortunately, most Statecharts variants do not possess a compositional semantics. *Compositionality*, however, is a necessary prerequisite for the re-usability of specifications and of their behavioral properties. The objective of the present work is to analyze the compositionality problem regarding the classical macro-step semantics of Statecharts, as introduced by Pnueli and Shalev. This analysis will provide the basis for deriving a compositional and *fully abstract* Statecharts semantics.

The compositionality defect of Pnueli and Shalev's macro-step semantics is illustrated by showing that equality of response behavior is not preserved by the concurrency and hierarchy operators of Statecharts. The reason is that this semantics abstracts from the causal reasoning employed when inferring macro steps, thereby imposing a closed-world assumption. Indeed, the compositionality problem can be traced back to the *invalidity of the Law of the Excluded Middle*. For the analysis, we interpreted Statecharts, relative to a given system state, as propositional formulas. These are given meaning as specific intuitionistic Kripke structures, namely linear chains encoding interactions between Statecharts and environments, which we refer to as stabilization sequences. On this semantic basis, we developed a compositional and fully abstract macro-step semantics in two steps. First, we restricted ourselves to studying Statecharts not containing any hierarchy operator. We showed that Pnueli and Shalev's semantics is contained in our novel semantics when considering non-extensible stabilization sequences of length one only. When also taking sequences of length two into account, which represent a single interaction with the environment, a compositional macro-step semantics is obtained. This semantics proves to be fully abstract regarding Pnueli and Shalev's semantics,

as it characterizes the largest congruence contained in equality on response behavior. Second, by presenting a normalization result that allows hierarchy operators to be pushed outside, we lifted our results to arbitrary Statecharts.

We will apply intuitionistic framework to uniformly comparing many variants of Statecharts semantics studied in the literature. The framework also provides a means for developing *algebraic characterizations* of macro-step semantics.

This work is done in collaboration with Michael Mendler (University of Sheffield, England).

A Novel Compositional Approach to Defining Statecharts Semantics

Gerald Lüttgen

Statecharts is a visual language for specifying reactive system behavior, which is founded on hierarchical, concurrent state machines. Despite its appeal to software engineers, its step semantics lacks the fundamental property of *compositionality*. Compositionality, however, is a prerequisite for exploiting the modular structure of Statecharts specifications for simulation, verification, and code generation. The aim of this research project is to promote a novel approach to defining Statecharts semantics which combines compositionality and the three semantic dimensions underlying Statecharts, namely causality, the synchrony hypothesis, and global consistency, while still being comprehensible.

The considered approach borrows from ideas investigated for timed process calculi which also employ the synchrony hypothesis and which allow one to represent ordinary system behavior and clock ticks using flat labeled transition systems. These transition systems are defined via *structural operational rules* along the state hierarchy of Statecharts. Our semantics explicitly represents macro steps as sequences of micro steps which begin and end with the ticking of a global clock. Thereby, compositionality is achieved on the explicit micro-step level and causality and synchrony on the implicit macro-step level. The present research originates with our work described in the Semiannual Report for October 1998–March 1999, which developed a compositional timed process algebra with priority. This algebra permits an embedding of a simple variant of Statecharts that is close to the classical variant considered by Pnueli and Shalev. Thereby, Statecharts is indirectly given a compositional operational semantics. In this project, we re-developed and explained our semantics without referring to process algebras, which eliminates the rather complicated indirection and makes the underlying semantic issues and design decisions more accessible to software engineers. Additionally, we proved the flexibility and elegance of our semantic framework by incorporating popular Statecharts features used in daily practice, such as state references, history states, and priority concepts.

We plan to employ our approach for interfacing various Statecharts dialects to *verification tools*. Moreover, we intend to adapt behavioral equivalences studied for process algebras to Statecharts.

This work is done in collaboration with Rance Cleaveland (SUNY at Stony Brook) and Michael von der Beeck (TU Munich, Germany).

Arcade: A Distributed Computing Environment for ICASE

Piyush Mehrotra

Distributed heterogeneous computing is being increasingly applied to a variety of large size computational problems. Such computations, for example, the multidisciplinary design optimization of an aircraft, generally consists of multiple heterogeneous modules interacting with each other to solve the problem at hand. These applications are generally developed by a team in which each discipline is the responsibility of experts in the field. The objective of this project is to develop a GUI based environment which supports the multi-user

design of such applications and their execution and monitoring in a heterogeneous environment consisting of a network of workstations, specialized machines and parallel architectures.

We have been implementing a Java-based three-tier prototype system which support a thin client interface for the design and execution of multi-module codes. The middle tier consists of logic to process the user input and also to manage the resource controllers which comprise the third tier. In the last six months we have continued our focus on the Jini based distributed resource monitoring and managing system. We had earlier extended Jini in order to use it across distributed domains. We have been testing and optimizing the implementation making changes and modifications where necessary. We are also implementing an XML based resource specification system which will ease the task of managing and using remote resources. We have also started the design of an XML-based project specification language in order to allow users to create Arcade project using scripts. We are integrating this in the current Arcade system building translators from the XML-based script to the internal project class and vice-versa.

In the future we will continue to work on the enhancements outlined above. In addition we will enhance the current Arcade monitoring system by adding more visualization tools. We will also add support for dynamic steering of executing simulation codes.

This work is being performed in collaboration with A. Al-Theneyan and M. Zubair (Old Dominion University).

Languages for High Performance and Distributed Computing

Piyush Mehrotra

There are many approaches to exploiting the power of parallel and distributed computers. Under this project, our focus is to evaluate these different approaches, proposing extensions and new compilation techniques where appropriate.

OpenMP is a set of directives extending C, C++ and Fortran which provide a shared-memory parallel programming model. Current parallel architectures are built by interconnecting nodes which internally provide true-shared memory across a small number of processors. Both hardware and software based approaches are used to provide a shared address space across the physically distributed memories of the nodes in larger systems. OpenMP provides an easy and incremental approach for small scale shared memory systems. However, controlling and exploiting data locality becomes an issue as the latency for data transfers across the nodes increases for larger systems. In the last few months we initiated a small forum consisting a group of researchers from academia, industry and the national labs interested in this topic. The forum had an initial meeting and is focussing on the issue of the requirements for locality control within OpenMP.

We plan to use the directions set by the forum to investigate distribution directives within OpenMP. This work is being done in collaboration with B. Chapman (University of Houston).

A Generic Higher-order Unification Algorithm

César A. Muñoz

Higher-Order Unification (HOU) is the process to solve equations in the λ-calculus. HOU problems arise in several program analysis techniques, for instance type checking, abstraction, optimization, and program transformation. HOU is also a main issue on advanced theorem provers and logic programming tools. Fundamental research on HOU are essential for the advancement of these efforts. The objective of this work is to investigate the application of a new technique called *explicit substitutions*, which allows substitutions to be first-class objects in λ-calculus, to formulate HOU-related problems for a general class of type systems.

It is well known that the λ-calculus can be encoded in a first-order setting via a nameless notation for variables called *de Bruijn indices*. First-class substitutions and de Bruijn indices notation allow a first-order presentation of HOU-related problems that eliminates technical nuisances due to the higher-order aspects of the λ-calculus. We have completed the meta-theoretical study of a λ-calculus of explicit substitutions for pure type systems. Using that framework, we have implemented a HOU algorithm for the Calculus of Constructions, a very rich type theory, in the functional object-oriented language Objective Caml. The soundness and completeness properties of the algorithm have been proved.

We are interested in extending our framework to internalize into the theory, external features of the λ-calculus as for example definitions and filling of place-holders. Furthermore, we plan to integrate the HOU algorithm to a constraint solver module for theorem provers and logic programming tools.

This work is done in collaboration with Nikolaj Bjørner (Kestrel Institute).

Verification of the AILS Alerting Algorithm

César A. Muñoz

The Airborne Information for Lateral Spacing (AILS) program at NASA Langley Research Center aims to give pilots the information necessary to make independent approaches to parallel runways with spacing down to 2500 feet in Instrument Meteorological Conditions. Independent parallel approaches to closely spaced runways allow to reduce traffic delays and increase airport efficiency. The AILS concept consists of accurate traffic information visible at the navigation display and an alerting algorithm which warns the crew when one of the aircraft involved in a parallel landing is diverting from the intended flight path. The objective of this work is to apply formal analysis and verification techniques to the AILS alerting algorithm to discover any possible errors, leading to unalerted near misses or collisions, that have not been detected during testing and simulation.

The AILS alerting algorithm analyzes the parallel aircraft states and makes time projections of possible intrusion scenarios. Based on these projections and risk criteria, the algorithm triggers a sequence of caution and warning alerts. In order to analyze the behavior of the algorithm, we have developed an abstract model of the algorithm and a model of parallel landing scenarios in the general verification system PVS. To visualize the general performance of the models, we also have implemented them in the programming language Java. Finally, we have formally verified, via the computer algebra tool MuPad, physical properties assumed as axioms in our formalization. Based on the formal models, we are currently verifying the correctness condition of the algorithm.

Once the formal analysis of the algorithm has been completed, we plan to apply our analysis methodology to study a new version of the algorithm that has been recently released to the public domain.

This work is done in collaboration with Victor Carreño (NASA Langley).

Scalable Parallel Algorithms for Incomplete Factorization Preconditioning

Alex Pothen

The parallel computation of robust preconditioners is a priority for solving large systems of equations iteratively in several scientific computing problems. We are developing parallel algorithms and software that can compute incomplete factorization preconditioners for high-level fill.

We assume that the adjacency graph of the coefficient matrix is well-partitionable: i.e., that it can be partitioned into subgraphs of roughly equal sizes such that each subgraph has few boundary nodes relative to the number of interior nodes. We map the subgraphs to processors, form a subdomain interconnection

graph, and order the subdomains by a graph coloring to reduce global dependences. On each subdomain, we locally reorder the interior vertices before the boundary vertices. This reordering limits the fill that joins a subgraph on one processor to a subgraph on another, and enhances the concurrency in the computation. The preconditioner computation takes places in two phases: in the first phase, each processor computes the rows of the preconditioner corresponding to the interior vertices of their subdomains. In the second phase, the rows corresponding to the boundary nodes are computed. This approach can make use of level-based and numerical threshold-based algorithms for computing preconditioners in parallel.

We have reported results on up to 216 processors of the SGI Origin, the Cray T3E, the Sun HPC 10000, and Coral, the ICASE Beowulf cluster. These results show that the algorithm is scalable, in agreement with our analysis of the algorithm. We are continuing to develop our parallel implementation.

This is joint work with David Hysom (Old Dominion University and ICASE).

Enhancing the Performance of Unstructured Mesh Codes

Alex Pothen

Irregular computations such as unstructured mesh algorithms or sparse matrix algorithms achieve only a small fraction of the available peak performance on modern microprocessors. Understanding the reasons for this low performance is difficult on modern microprocessors since they issue multiple instructions in each clock cycle, and permit them to execute out of issue-order by exploiting multiple functional units and multiple levels of caches. We have initiated a study measuring the performance of these codes using hardware performance counters and simulation tools to first understand and then improve the performance.

We have studied the performance of the kernel of an unstructured mesh code for solving Euler's equations. By measuring various performance metrics, we find that on an SGI architecture the performance is limited primarily by the number of loads and stores that can be issued in a clock cycle. By scheduling the code such that it makes effective use of the registers and functional units, we have been able to improve the performance of the kernel from about 40 Mflops/second to 270 Mflops/second, half the peak performance of a 270 MHz MIPS R12000 processor that can graduate two floating point instructions per clock cycle. One surprise so far has been that for this code, the level 1 and level 2 caches play a relatively minor role in determining performance on the SGI.

We are continuing our work to find general principles for improving the performance of irregular computations.

This is joint work with Jinghua Fu (Old Dominion University) and Dimitri Mavriplis (ICASE).

External Memory Sparse Direct Solvers

Alex Pothen

We have begun to develop external memory algorithms and software for solving large, sparse systems of equations by means of direct solvers. Such methods will enable sparse direct solvers to make effective use of the Terabytes of storage space available on today's serial computers, without having to employ parallel computers in order to meet the storage requirements of direct solvers.

Our work is part of a broader project that examines the design of sparse direct solvers that make effective use of multiple levels of the memory hierarchy from disks to registers. We have designed blocked algorithms that have optimal data movement complexity for the multiple levels; our implementation shows that these algorithms make effective use of each individual level in the memory hierarchy, and that the performance at each level is an upper bound on the performance that can be achieved at the next farther level from the processor.

We are currently characterizing algorithms that require the least time when a specified, fixed amount of primary memory is available, and schedules for computation that achieve these times.

This is joint work with Florin Dobrian (Old Dominion University).

Shop Scheduling with Application to Parallel Computing

Alex Povitsky

The need for due-date scheduling in networks, parallel numerical algorithms, high-performance computer graphics, operations research, and industrial engineering has been recognized for a while. Researchers distinguish between two different arrangements of the scheduling problem: (i) due dates are constraints of the problem; and (ii) deviations from due dates must be minimized. Most available studies consider the latter problem statement, whereas studies about the former formulation are restricted to the one-machine case with non-preemptive restriction and without precedence constraints between jobs. In the current study, we consider a problem of scheduling of a shop of m machines that process a set of n jobs where job due dates are given constraints (absolute deadlines). Our algorithm of processor scheduling for implicit finite-difference parallel numerical algorithms represents a special case of shop scheduling where jobs (groups of lines of numerical grids) are executed in either increasing or decreasing order of machines (processors) and treatment times are equal for all jobs and machines. The current study generalizes the aforementioned scheduling algorithm because each job has a different order of sequential operations (i.e., path of a job through the set of machines) and the number of time units per operation is different for various jobs and machines. If the goal is to find the absolute minimum of workspan, the job shop problem (i) becomes $NP-$hard and is recognized among the hardest combinatorial optimization problems.

Absolute deadlines cannot reduce the problem complexity and the job shop problem with absolute deadlines remains $NP-$hard. However, the given additional constraints on order of execution of jobs on machines and/or completion times lead to a considerable reduction of number of possible schedules. To minimize the shop workspan, we propose heuristic scheduling algorithms that are based on backward scheduling starting from the last outermost machine, where the completion times are given. In other words, we use a greedy algorithm to define schedule on all other machines starting from the last outermost machine. In some cases, we treat job shop problems as sequences of forward and backward flow shop problems. For the dynamic insertion of a new job with restriction on its completion and starting times, we propose a combination of backward scheduling with look ahead for the new job and forward rescheduling of tails of some existing jobs. This rescheduling is arranged in such a way that the maximum delay of jobs reaches its minimum. Our computer experiments show that for statistically balanced machines, where the overall load per machine is the sum of random functions of the same range, the proposed algorithm computes schedules with workspans that are close to their lower bound.

Our future research includes combinations of greedy algorithms with shifting bottleneck procedures. We shall implement the proposed methodology for numerical problems with an uneven multi-domain workload, for problems of design optimization and for scheduling of data-dependent multi-task problems.

Solving Radiation Transport Equations with the Multigrid Method

Linda Stals

We are working on an ASCI II project which is concerned with the solution of radiation transport equations. The solution of such equations is very computationally expensive, pushing the limits of the resources which are currently available. The project is designed to investigate the use of algorithms which scale well as the problem size is increased currently focusing on a parallel multigrid algorithm.

We have investigated the case where all temperatures are in equilibrium and the model reduces to a single equation. This equation contains many of the features inherent in the full model, such as strong nonlinearities and large jumps in the coefficients. The equation was solved using both a Newton-Multigrid method and a Full Approximation Scheme (nonlinear multigrid) method. We found that the Newton-Multigrid method worked better because it did not need to calculate the diffusion coefficient as often. Such an observation is of interest because in more complex models the calculation of the diffusivity term will become even more expensive. This gives us a better idea of what solution methods to try when looking at systems of equations.

In future work we want to include more physics in the equations that we are solving. The PI's for the ASCI II project are David Keyes and Alex Pothen (Old Dominion University and ICASE) and Dimitri Mavriplis (ICASE).

REPORTS AND ABSTRACTS

Rubinstein, Robert, and Ye Zhou: *Characterization of sound radiation by unresolved scales of motion in computational aeroacoustics.* ICASE Report No. 99-39, (NASA/CR-1999-209688), October 13, 1999, 13 pages. Submitted to the European Journal of Mechanics B.

Evaluation of the sound sources in a high Reynolds number turbulent flow requires time-accurate resolution of an extremely large number of scales of motion. Direct numerical simulations will therefore remain infeasible for the forseeable future; although current large eddy simulation methods can resolve the largest scales of motion accurately, they must leave some scales of motion unresolved. A priori studies show that acoustic power can be underestimated significantly if the contribution of these unresolved scales is simply neglected.

In this paper, the problem of evaluating the sound radiation properties of the unresolved, subgrid-scale motions is approached in the spirit of the simplest subgrid stress models: the unresolved velocity field is treated as isotropic turbulence with statistical descriptors evaluated from the resolved field. The theory of isotropic turbulence is applied to derive formulas for the total power and the power spectral density of the sound radiated by a filtered velocity field. These quantities are compared with the corresponding quantities for the unfiltered field for a range of filter widths and Reynolds numbers.

Povitsky, Alex: *Wavefront cache-friendly algorithm for compact numerical schemes.* ICASE Report No. 99-40, (NASA/CR-1999-209708), October 15, 1999, 12 pages. Submitted to Applied Mathematics Letters.

Compact numerical schemes provide high-order solution of PDEs with low dissipation and dispersion. Computer implementation of these schemes requires numerous passes of data through cache memory that considerably reduces performance of these schemes. To reduce this difficulty, a novel algorithm is proposed here. This algorithm is based on a wavefront approach and sweeps through cache only twice.

Ma, Kwan-Liu, and Thomas W. Crockett: *Parallel visualization of large-scale aerodynamics calculations: A case study on the Cray T3E.* ICASE Report No. 99-41, (NASA/CR-1999-209709), October 18, 1999, 17 pages. To appear in the Proceedings of the 1999 Symposium on Parallel Visualization and Graphics.

This paper reports the performance of a parallel volume rendering algorithm for visualizing a large-scale unstructured-grid dataset produced by a three-dimensional aerodynamics simulation. This dataset, containing over 18 million tetrahedra, allows us to extend our performance results to a problem which is more than 30 times larger than the one we examined previously. This high resolution dataset also allows us to see fine, three-dimensional features in the flow field. All our tests were performed on the SGI/Cray T3E operated by NASA's Goddard Space Flight Center. Using 511 processors, a rendering rate of almost 9 million tetrahedra/second was achieved with a parallel overhead of 26%.

Lüttgen, Gerald, Michael von der Beeck, and Rance Cleaveland: *Statecharts via process algebra.* ICASE Report No. 99-42, (NASA/CR-1999-209713), October 26, 1999, 23 pages. Presented at the International Conference on Concurrency Theory (CONCUR'99).

Statecharts is a visual language for specifying the behavior of reactive systems. The language extends finite-state machines with concepts of hierarchy, concurrency, and priority. Despite its popularity as a design

notation for embedded systems, precisely defining its semantics has proved extremely challenging. In this paper, a simple process algebra, called Statecharts Process Language (SPL), is presented, which is expressive enough for encoding Statecharts in a structure-preserving and semantics-preserving manner. It is established that the behavioral relation bisimulation, when applied to SPL, preserves Statecharts semantics.

Muñoz, César: *Dependent types and explicit substitutions.* ICASE Report No. 99-43, (NASA/CR-1999-209722), November 5, 1999, 31 pages. Submitted to Mathematical Structures in Computer Science.

We present a dependent-type system for a λ-calculus with explicit substitutions. In this system, meta-variables, as well as substitutions, are first-class objects. We show that the system enjoys properties like type uniqueness, subject reduction, soundness, confluence and weak normalization.

Mavriplis, Dimitri J.: *Large-scale parallel viscous flow computations using an unstructured multigrid algorithm.* ICASE Report No. 99-44, (NASA/CR-1999-209724), November 10, 1999, 19 pages. To appear in the Proceedings for Parallel CFD'99.

The development and testing of a parallel unstructured agglomeration multigrid algorithm for steady-state aerodynamic flows is discussed. The agglomeration multigrid strategy uses a graph algorithm to construct the coarse multigrid levels from the given fine grid, similar to an algebraic multigrid approach, but operates directly on the non-linear system using the FAS approach. The scalability and convergence rate of the multigrid algorithm are examined on the SGI Origin 2000 and the Cray T3E. An argument is given which indicates that the asymptotic scalability of the multigrid algorithm should be similar to that of its underlying single grid smoothing scheme. For medium size problems involving several million grid points, near perfect scalability is obtained for the single grid algorithm, while only a slight drop-off in parallel efficiency is observed for the multigrid V- and W-cycles, using up to 128 processors on the SGI Origin 2000, and up to 512 processors on the Cray T3E. For a large problem using 25 million grid points, good scalability is observed for the multigrid algorithm using up to 1450 processors on a Cray T3E, even when the coarsest grid level contains fewer points than the total number of processors.

Woodruff, S.L., J.V. Shebalin, and M.Y. Hussaini: *Direct-numerical and large-eddy simulations of a non-equilibrium turbulent Kolmogorov flow.* ICASE Report No. 99-45, (NASA/CR-1999-209727), January 7, 2000, 21 pages. To be submitted to Physics of Fluids.

A non-equilibrium form of turbulent Kolmogorov flow is set up by making an instantaneous change in the amplitude of the spatially- periodic forcing. It is found that the response of the flow to this instantaneous change becomes more dramatic as the wavenumber of the forcing is increased, and, at the same time, that the faithfulness with which the large-eddy-simulation results agree with the direct-numerical results decreases.

Yamaleev, Nail K.: *Minimization of the truncation error by grid adaptation.* ICASE Report No. 99-46, (NASA/CR-1999-209729), December 2, 1999, 40 pages. To be submitted to the Journal of Computational Physics.

A new grid adaptation strategy, which minimizes the truncation error of a pth-order finite difference approximation, is proposed. The main idea of the method is based on the observation that the local truncation error associated with discretization on nonuniform meshes can be minimized if the interior grid points

are redistributed in an optimal sequence. The method does not explicitly require the truncation error estimate and at the same time, it allows one to increase the design order of approximation by one globally, so that the same finite difference operator reveals superconvergence properties on the optimal grid. Another very important characteristic of the method is that if the differential operator and the metric coefficients are evaluated identically by some hybrid approximation the single optimal grid generator can be employed in the entire computational domain independently of points where the hybrid discretization switches from one approximation to another. Generalization of the present method to multiple dimensions is presented. Numerical calculations of several one-dimensional and one two-dimensional test examples demonstrate the performance of the method and corroborate the theoretical results.

Muñoz, César: *Proof-term synthesis on dependent-type systems via explicit substitutions.* ICASE Report No. 99-47, (NASA/CR-1999-209730), November 30, 1999, 33 pages. Submitted to Theoretical Computer Science.

Typed λ-terms are used as a compact and linear representation of proofs in intuitionistic logic. This is possible since the Curry-Howard isomorphism relates proof trees with typed λ-terms. The proofs-as-terms principle can be used to check a proof by type checking the λ-term extracted from the complete proof tree. However, proof trees and typed λ-terms are built differently. Usually, an auxiliary representation of unfinished proofs is needed, where type checking is possible only on complete proofs. In this paper we present a proof synthesis method for dependent-type systems where typed open terms are built incrementally at the same time as proofs are done. This way, every construction step, not just the last one, may be type checked. The method is based on a suitable calculus where substitutions as well as meta-variables are first-class objects.

Povitsky, Alex: *On aeroacoustics of a stagnation flow near a rigid wall.* ICASE Report No. 99-48, (NASA/CR-1999-209825), December 6, 1999, 23 pages. To be submitted to Physics of Fluids.

We consider propagation of disturbances in a non-uniform mean flow by high-order numerical simulation. Monopole and dipole acoustic, vortical and entropy pulses are embedded in an incompressible stagnation flow, which is taken as a prototype of a non-uniform low Mach number mean flow near a rigid wall at high angle of attack.

Numerical results are discussed in terms of baroclinic generation of disturbance vorticity that appear to be a key process in energy transfer between a non-uniform mean flow and a propagating disturbance. These phenomena lead to amplification of sound waves originated from an acoustic pulse. Vorticity generation governs wave radiation of a near-wall entropy pulse and makes the radiated waves similar to those from a vortical dipole. Interaction of initial pulse vorticity with generated vorticity leads to various radiated wave patterns discussed here.

Alexandrov, Natalia M., Robert Michael Lewis, Clyde R. Gumbert, Larry L. Green, and Perry A. Newman: *Optimization with variable-fidelity models applied to wing design.* ICASE Report No. 99-49, (NASA/CR-1999-209826), December 6, 1999, 23 pages. AIAA Paper 2000-0841.

This work discusses an approach, the Approximation Management Framework (AMF), for solving optimization problems that involve computationally expensive simulations. AMF aims to maximize the use of lower-fidelity, cheaper models in iterative procedures with occasional, but systematic, recourse to higher-fidelity, more expensive models for monitoring the progress of the algorithm. The method is globally convergent to a solution of the original, high-fidelity problem. Three versions of AMF, based on three nonlinear

programming algorithms, are demonstrated on a 3D aerodynamic wing optimization problem and a 2D airfoil optimization problem. In both cases Euler analysis solved on meshes of various refinement provides a suite of variable-fidelity models. Preliminary results indicate threefold savings in terms of high-fidelity analyses in case of the 3D problem and twofold savings for the 2D problem.

Ciardo, Gianfranco, Gerald Lüttgen, and Radu Siminiceanu: *Efficient symbolic state-space construction for asynchronous systems.* ICASE Report No. 99-50, (NASA/CR-1999-209827), December 13, 1999, 43 pages. Submitted to the 21st International Conference on Applications and Theory on Petri Nets.

Many state-of-the-art techniques for the verification of today's complex embedded systems rely on the analysis of their reachable state spaces. In this paper, we develop a new algorithm for the symbolic generation of the state space of asynchronous system models, such as Petri nets. The algorithm is based on previous work that employs Multi-valued Decision Diagrams (MDDs) for efficiently storing sets of reachable states. In contrast to related approaches, however, it fully exploits event locality which is a fundamental semantic property of asynchronous systems. Additionally, the algorithm supports intelligent cache management and achieves faster convergence via advanced iteration control. It is implemented in the tool SMART, and run-time results for several examples taken from the Petri net literature show that the algorithm performs about one order of magnitude faster than the best existing state-space generators.

Thomas, James L., Boris Diskin, and Achi Brandt: *Textbook multigrid efficiency for the incompressible Navier-Stokes equations: High Reynolds number wakes and boundary layers.* ICASE Report No. 99-51, (NASA/CR-1999-209831), December 22, 1999, 27 pages. To be submitted to Computers and Fluids.

Textbook multigrid efficiencies for high Reynolds number simulations based on the incompressible Navier-Stokes equations are attained for a model problem of flow past a finite flat plate. Elements of the Full Approximation Scheme multigrid algorithm, including distributed relaxation, defect correction, and boundary treatment, are presented for the three main physical aspects encountered: entering flow, wake flow, and boundary layer flow. Textbook efficiencies, i.e., reduction of algebraic errors below discretization errors in one full multigrid cycle, are attained for second order accurate simulations at a laminar Reynolds number of 10,000.

Smith, Ralph C., and Zoubeida Ounaies: *A domain wall model for hysteresis in piezoelectric materials.* ICASE Report No. 99-52, (NASA/CR-1999-209832), December 22, 1999, 21 pages. Submitted to the Journal of Intelligent Material Systems and Structures.

This paper addresses the modeling of hysteresis and nonlinear constitutive relations in piezoelectric materials at moderate to high drive levels. Hysteresis and nonlinearities are due to the domain structure inherent to the materials and both aspects must be addressed to attain the full potential of the materials as sensors and actuators in high performance applications. The model employed here is based on previously developed theory for hysteresis in general ferroelectric materials. This theory is based on the quantification of the reversible and irreversible motion of domain walls pinned at inclusions in the material. This yields an ODE model having five parameters. The relationship of the parameters to physical attributes of the materials is detailed and algorithms for determining estimates of the parameters using measured values of the coercive field, differential susceptibility and saturation properties of the materials are detailed. The accuracy of the model and its capability for the prediction of measured polarization at various drive levels

is illustrated through a comparison with experimental data from PZT5A, PZT5H and PZT4 compounds. Finally, the ODE model formulation is amenable to inversion which facilitates the construction of an inverse compensator for linear control design.

Park, C., Z. Ounaies, J. Su, J.G. Smith, Jr., and J.S. Harrison: *Polarization stability of amorphous piezoelectric polyimides.* ICASE Report No. 99-53, (NASA/CR-1999-209833), December 22, 1999, 11 pages. To appear in the Proceedings of MRS '99.

Amorphous polyimides containing polar functional groups have been synthesized and investigated for potential use as high temperature piezoelectric sensors. The thermal stability of the piezoelectric effect of one polyimide was evaluated as a function of various curing and poling conditions under dynamic and static thermal stimuli. First, the polymer samples were thermally cycled under strain by systematically increasing the maximum temperature from 50°C to 200°C while the piezoelectric strain coefficient was being measured. Second, the samples were isothermally aged at an elevated temperature in air, and the isothermal decay of the remanent polarization was measured at room temperature as a function of time. Both conventional and corona poling methods were evaluated. This material exhibited good thermal stability of the piezoelectric properties up to 100°C.

Nordström, Jan, and Mark H. Carpenter: *High order finite difference methods, multidimensional linear problems and curvilinear coordinates.* ICASE Report No. 99-54, (NASA/CR-1999-209834), January 7, 2000, 35 pages. To be submitted to the Journal of Computational Physics.

Boundary and interface conditions are derived for high order finite difference methods applied to multidimensional linear problems in curvilinear coordinates. The boundary and interface conditions lead to conservative schemes and strict and strong stability provided that certain metric conditions are met.

Smith, Ralph C., and Zoubeida Ounaies: *A model for assymetric hysteresis in piezoceramic materials.* ICASE Report No. 2000-1, (NASA/CR-2000-209837), January 13, 2000, 11 pages. To appear in the Proceedings of the 1999 Fall Meeting of the Material Research Society, Boston, MA.

This paper focuses on the characterization of hysteresis exhibited by piezoelectric materials at moderate to high field levels. For soft materials in which dipoles are easily reconfigured, the hysteresis loop is observed to be rotationally symmetric about the zero field, zero polarization point and symmetric models can be employed. In harder materials, however, the loops are no longer rotationally symmetric which necessitates the development of commensurate characterization techniques. The model considered here is based upon the quantification of reversible and irreversible changes in polarization due to bending and translation of domain walls pinned at inclusions inherent to the materials. The performance of the model is illustrated through comparison with PZT4 data.

Anderson, W.K., W.D. Gropp, D.K. Kaushik, D.E. Keyes, and B.F. Smith: *Achieving high sustained performance in an unstructured mesh CFD application.* ICASE Report No. 2000-2, (NASA/CR-2000-209838), January 31, 2000, 16 pages. To appear in the Proceedings of SC'99.

This paper highlights a three-year project by an interdisciplinary team on a legacy F77 computational fluid dynamics code, with the aim of demonstrating that implicit unstructured grid simulations can execute at

rates not far from those of explicit structured grid codes, provided attention is paid to data motion complexity and the reuse of data positioned at the levels of the memory hierarchy closest to the processor, in addition to traditional operation count complexity. The demonstration code is from NASA and the enabling parallel hardware and (freely available) software toolkit are from DOE, but the resulting methodology should be broadly applicable, and the hardware limitations exposed should allow programmers and vendors of parallel platforms to focus with greater encouragement on sparse codes with indirect addressing. This snapshot of ongoing work shows a performance of 15 microseconds per degree of freedom to steady-state convergence of Euler flow on a mesh with 2.8 million vertices using 3072 dual-processor nodes of Sandia's "ASCI Red" Intel machine, corresponding to a sustained floating-point rate of 0.227 Tflop/s.

Rubinstein, Robert, and Aaron H. Auslender: *Non-ideal gas effects on shock waves in weakly ionized gases.* ICASE Report No. 2000-3, (NASA/CR-2000-209839), January 31, 2000, 13 pages. To be submitted to Physical Review E.

A shock wave in a weakly ionized gas can be preceded by a charge separation region if the Debye length is larger than the shock width. It has been proposed that electrostatic contributions to pressure in the charge separation region can increase the sound speed ahead of the shock well above the sound speed in a neutral gas at the same temperature and therefore increase the shock propagation speed. This proposal is investigated numerically and theoretically. It is concluded that although the ion gas becomes strongly non-ideal in the charge separation region, there is no appreciable effect on the neutral shock.

Lüttgen, Gerald, César Muñoz, Ricky Butler, Ben Di Vito, and Paul Miner: *Towards a customizable PVS.* ICASE Report No. 2000-4, (NASA/CR-2000-209851), February 2, 2000, 19 pages. Presented at the PVS User Group Meeting, World Congress on Formal Methods (FM'99).

PVS is a state-of-the-art theorem-proving tool developed by SRI International. It is used in a variety of academic and real-world applications by NASA and ICASE researchers, for whom tool customization and extensibility are becoming increasingly important issues. This paper shows, by referring to past experiences with several projects and case studies, that the customization features currently offered by PVS are often insufficient. It also suggests several improvements regarding PVS's customization in the short run and regarding its extensibility in the long run.

Rubinstein, Robert: *Formulation of a two-scale model of turbulence.* ICASE Report No. 2000-5, (NASA/CR-2000-209853), February 16, 2000, 18 pages. To be submitted to Physics of Fluids.

A two-scale turbulence model is derived by averaging the two-point spectral evolution equation. In this model, the inertial range energy transfer and the dissipation rate can be unequal. The model is shown to reduce to a standard two-equation model in decaying turbulence.

Mei, Renwei, Li-Shi Luo, and Wei Shyy: *An accurate curved boundary treatment in the lattice Boltzmann method.* ICASE Report No. 2000-6, (NASA/CR-2000-209854), February 16, 2000, 25 pages. Journal of Computational Physics, Vol. 115, 1999, pp. 307–330.

The lattice Boltzmann equation (LBE) is an alternative kinetic method capable of solving hydrodynamics for various systems. Major advantages of the method are owing to the fact that the solution for the particle

distribution functions is explicit, easy to implement, and natural to parallelize. Because the method often uses uniform regular Cartesian lattices in space, curved boundaries are often approximated by a series of stairs that leads to reduction in computational accuracy. In this work, a second-order accurate treatment of boundary condition in the LBE method is developed for a curved boundary. The proposed treatment of the curved boundaries is an improvement of a scheme due to Filippova and Hanel. The proposed treatment for curved boundaries is tested against several flow problems: 2-D channel flows with constant and oscillating pressure gradients for which analytic solutions are known, flow due to an impulsively started wall, lid-driven square cavity flow, and uniform flow over a column of circular cylinders. The second-order accuracy is observed with solid boundary arbitrarily placed between lattice nodes. The proposed boundary condition has well behaved stability characteristics when the relaxation time is close to 1/2, the zero limit of viscosity. The improvement can make a substantial contribution toward simulating practical fluid flow problems using the lattice Boltzmann method.

Luo, Li-Shi: *Some recent results on discrete velocity model and ramifications for lattice Boltzmann equation.* ICASE Report No. 2000-7, (NASA/CR-2000-209855), February 16, 2000, 18 pages. To appear in Computer Physics Communications.

Some rigorous results on discrete velocity models are briefly reviewed and their ramifications for the lattice Boltzmann equation (LBE) are discussed. In particular, issues related to thermodynamics and H-theorem of the lattice Boltzmann equation are addressed. It is argued that for the lattice Boltzmann equation satisfying the correct hydrodynamic equations, there cannot exist an H-theorem. Nevertheless, the equilibrium distribution function of the lattice Boltzmann equation can closely approximate the genuine equilibrium which minimizes the H-function of the corresponding continuous Boltzmann equation. It is also pointed out that the "equilibrium" in the LBE models is an attractor rather than a true equilibrium in the rigorous sense of H-theorem. Since there is no H-theorem to guarantee the stability of the LBE models at the attractor, the stability of the attractor can only be studied by means other than proving an H-function.

Jordan, T., Z. Ounaies, J. Tripp, and P. Tcheng: *Electrical properties and power considerations of a piezoelectric actuator.* ICASE Report No. 2000-8, (NASA/CR-2000-209861), February 18, 2000, 13 pages. To appear in the Proceedings of MRS '99.

This paper assesses the electrical characteristics of piezoelectric wafers for use in aeronautical applications such as active noise control in aircraft. Determination of capacitive behavior and power consumption is necessary to optimize the system configuration and to design efficient driving electronics. Empirical relations are developed from experimental data to predict the capacitance and loss tangent of a PZT5A ceramic as nonlinear functions of both applied peak voltage and driving frequency. Power consumed by the PZT is the rate of energy required to excite the piezoelectric system along with power dissipated due to dielectric loss and mechanical and structural damping. Overall power consumption is thus quantified as a function of peak applied voltage and driving frequency. It was demonstrated that by incorporating the variation of capacitance and power loss with voltage and frequency, satisfactory estimates of power requirements can be obtained. These relations allow general guidelines in selection and application of piezoelectric actuators and driving electronics for active control applications.

Lončarić, J., V.S. Ryaben'kii, and S.V. Tsynkov: *Active shielding and control of environmental noise.* ICASE Report No. 2000-9, (NASA/CR-2000-209862), February 18, 2000, 35 pages. To be submitted to the SIAM Journal of Control and Optimization.

We present a mathematical framework for the active control of time-harmonic acoustic disturbances. Unlike many existing methodologies, our approach provides for the exact volumetric cancellation of unwanted noise in a given predetermined region of space while leaving unaltered those components of the total acoustic field that aredeemed as friendly. Our key finding is that for eliminating the unwanted component of the acoustic field in a given area, one needs to know relatively little; in particular, neither the locations nor structure nor strength of the exterior noise sources needs to be known. Likewise, there is no need to know the volumetric properties of the supporting medium across which the acoustic signals propagate, except, maybe, in the narrow area of space near the boundary (perimeter) of the domain to be shielded. The controls are built based solely on the measurements performed on the perimeter of the region to be shielded; moreover, the controls themselves (i.e., additional sources) are concentrated also only near this perimeter. Perhaps as important, the measured quantities can refer to the total acoustic field rather than to its unwanted component only, and the methodology can automatically distinguish between the two.

In the paper, we construct a general solution to the aforementioned noise control problem. The apparatus used for deriving the general solution is closely connected to the concepts of generalized potentials and boundary projections of Calderon's type. For a given total wave field, the application of Calderon's projections allows us to definitively split between its incoming and outgoing components with respect to a particular domain of interest, which may have arbitrary shape. Then, the controls are designed so that they suppress the incoming component for the domain to be shielded or alternatively, the outgoing component for the domain, which is complementary to the one to be shielded. To demonstrate that the new noise control technique is appropriate, we thoroughly work out a two-dimensional model example that allows full analytical consideration. To conclude, we very briefly discuss the numerical (finite-difference) framework for active noise control that has, in fact, already been worked out, as well as some forthcoming extensions of the current work: Optimization of the solution according to different criteria that would fit different practical requirements, applicability of the new technique to quasi-stationary problems, and active shielding in the case of the broad-band spectra of disturbances.

Joyner, Michele L., H.T. Banks, Buzz Wincheski, and William P. Winfree: *Nondestructive evaluation using a reduced order computational methodology.* ICASE Report No. 2000-10, (NASA/CR-2000-209870), March 7, 2000, 22 pages. Submitted to Inverse Problems.

This paper uses eddy current based techniques and reduced order modeling to explore the feasibility of detecting a subsurface damage in structures such as air foils and pipelines. To identify the geometry of a damage, an optimization algorithm is employed which requires solving the forward problem numerous times. To implement these methods in a practical setting, the forward algorithm must be solved with extremely fast and accurate solution methods. Therefore, our computational methods are based on the reduced order Karhunen-Loeve or Proper Orthogonal Decomposition (POD) techniques. For proof-of-concept, we implement the methodology on a 2-D problem and find the methods to be efficient and robust even with data containing 10% relative noise. Furthermore, the methods are fast; our findings suggest we can reduce the computational time on average by a factor of 3000.

Nicol, David M., and Gianfranco Ciardo: *Automated parallelization of discrete state-space generation.* ICASE Report No. 2000-11, (NASA/CR-2000-210082), March 28, 2000, 24 pages. Journal of Parallel and Distributed Computing, Vol. 47, 1997, pp. 153-167.

We consider the problem of generating a large state-space in a distributed fashion. Unlike previously proposed solutions that partition the set of reachable states according to a hashing function provided by the user, we explore heuristic methods that completely automate the process. The first step is an initial random walk through the state space to initialize a search tree, duplicated in each processor. Then, the reachability graph is built in a distributed way, using the search tree to assign each newly found state to classes assigned to the available processors. Furthermore, we explore two remapping criteria that attempt to balance memory usage or future workload, respectively. We show how the cost of computing the global snapshot required for remapping will scale up for system sizes in the foreseeable future. An extensive set of results is presented to support our conclusions that remapping is extremely beneficial.

Lüttgen, Gerald, Michael von der Beeck, and Rance Cleaveland: *A compositional approach to Statecharts semantics.* ICASE Report No. 2000-12, (NASA/CR-2000-210086), March 28, 2000, 20 pages. Submitted to ACM SIGSOFT 2000 Eighth International Symposium on Foundations of Software Engineering (FSE-8).

Statecharts is a visual language for specifying reactive system behavior. The formalism extends traditional finite-state machines with notions of hierarchy and concurrency, and it is used in many popular software design notations. A large part of the appeal of Statecharts derives from its basis in state machines, with their intuitive operational interpretation. The traditional semantics of Statecharts, however, suffers from a serious defect: it is not compositional, meaning that the behavior of system descriptions cannot be inferred from the behavior of their subsystems. Compositionality is a prerequisite for exploiting the modular structure of Statecharts for simulation, verification, and code generation, and it also provides the necessary foundation for reusability.

This paper suggests a new compositional approach to formalizing Statecharts semantics as flattened transition systems in which transitions represent system steps. The approach builds on ideas developed for timed process calculi and employs structural operational rules to define the transitions of a Statecharts expression in terms of the transitions of its subexpressions. It is first investigated for a simple dialect of Statecharts, with respect to a variant of Pnueli and Shalev's semantics, and is illustrated by means of a small example. To demonstrate its flexibility, the proposed approach is then extended to deal with practically useful features available in many Statecharts variants, namely state references, history states, and priority concepts along state hierarchies.

Mavriplis, Dimitri J.: *Parallel investigations of an unstructured mesh Navier-Stokes solver.* ICASE Report No. 2000-13, (NASA/CR-2000-210088), March 31, 2000, 23 pages.

A Reynolds-averaged Navier-Stokes solver based on unstructured mesh techniques for analysis of high-lift configurations is described. The method makes use of an agglomeration multigrid solver for convergence acceleration. Implicit line-smoothing is employed to relieve the stiffness associated with highly stretched meshes. A GMRES technique is also implemented to speed convergence at the expense of additional memory usage. The solver is cache efficient and fully vectorizable, and is parallelized using a two-level hybrid MPI-OpenMP implementation suitable for shared and/or distributed memory architectures, as well as clusters of shared memory machines. Convergence and scalability results are illustrated for various high-lift cases.

Cleaveland, Rance, and Gerald Lüttgen: *Model checking is refinement — Relating Büchi testing and linear-time temporal logic —*. ICASE Report No. 2000-14, (NASA/CR-2000-210090), March 31, 2000, 25 pages. Submitted to the 25th International Symposium on Mathematical Foundations of Computer Science.

This paper develops a semantic foundation for reasoning about reactive systems specifications featuring combinations of labeled transition systems and formulas in linear-time temporal logic (LTL). Using Buechi automata as a semantic basis, the paper introduces two refinement preorders based on DeNicola and Hennessy's notion of may- and must-testing. Alternative characterizations for these relations are provided and used to show that the new preorders are conservative extensions of the traditional DeNicola and Hennessy preorders. The paper then establishes a tight connection between LTL formula satisfaction and the Buechi must-preorder. More precisely, it is shown that a labeled transition system satisfies an LTL formula if and only if it refines an appropriately defined Buechi automaton that can be constructed from the formula. Consequently, the Buechi must-preorder allows for a uniform treatment of traditional notions of process refinement and model checking. The implications of the novel theory are illustrated by means of a simple example system, in which some components are specified as transition systems and others as LTL formulas.

ICASE COLLOQUIA

Name/Affiliation/Title	Date
Schutzenhofer, Luke, University of Alabama, Huntsville "Launch Vehicle Design Process: Characterization, Technical Integration, and Lessons Learned"	October 8
Wainwright, Steven, Duke University "Biomimetics Seminar: Swimming by Design: Fish, Whales, and Submarines"	November 2
Klarlund, Nils, AT&T Labs Research "Mona Tutorial: Automaton-based Symbolic Computation"	November 4
Nelson, Michael L., NASA Langley Research Center "Smart Objects and Dumb Archives: Using Buckets in Digital Libraries"	November 9
Smalley, Richard E., Rice University "Buckytubes – New Materials and New Devices from Carbon"	November 19
Dowell, Earl, Duke University "Reduced Order Models in Unsteady Aerodynamics"	December 7
Camassa, Roberto, University of North Carolina, Chapel Hill "Fully Nonlinear Models for Internal Wave Propagation in Two-fluid Systems"	December 10
Yu, Dongsheng, McGill University "Dendritic Crystal Growth: The Effects of Fluid Flow"	January 7
Lee, Dong-Chan, Wichita State University "Momentum Vector Control for Spin Recovery"	January 24
Sahay, Anupam, University of Wisconsin, Madison "Models for Wall-bounded Shear Flows"	February 25
Bower, Christopher, University of North Carolina, Chapel Hill "Carbon Nanotubes: Fabrication, Properties, and Applications"	February 28
Rasmussen, Steen, Los Alamos National Laboratory ICASE Series on Computational Nanotechnology: "Molecular Dynamics Lattice Gas"	March 2
Tew, Gregory, University of Illinois, Urbana-Champaign "Assembling Molecules into Ordered Materials"	March 3

Name/Affiliation/Title	Date
Cha, Jennifer, University of California, Santa Barbara "Lessons from Nature: Novel Routes for Biomimetic Synthesis of Ordered Silica Structures"	March 6
He, Guowei, Los Alamos National Laboratory "Statistics of Small Scale Turbulence"	March 9
Haverkort, Boudewijn, RWTH Aachen, Germany "Towards Analysing Large Markov Chains Using Model Checking"	March 23
Boggavarapu, Sajiv, University of Arizona "Biomimetic Materials: From Artificial Antibodies to Controlled Crystal Growth?"	March 27

ICASE STAFF

I. ADMINISTRATIVE

Manuel D. Salas, Director, M.S., Aeronautics and Astronautics, Polytechnic Institute of Brooklyn, 1970. Fluid Mechanics and Numerical Analysis.

Linda T. Johnson, Office and Financial Administrator

Barbara A. Cardasis, Administrative Secretary

Etta M. Morgan, Accounting Supervisor

Emily N. Todd, Conference Manager/Executive Assistant

Shannon K. Verstynen, Information Technologist

Gwendolyn W. Wesson, Contract Accounting Clerk

Shouben Zhou, Systems Manager

Peter J. Kearney, Student Assistant

II. SCIENCE COUNCIL

David Gottlieb, (Chair) Professor, Division of Applied Mathematics, Brown University.

Lee Beach, Professor, Department of Physics, Computer Science & Engineering, Christopher Newport University.

Joseph E. Flaherty, Amos Eaton Professor, Departments of Computer Science and Mathematical Sciences, Rensselaer Polytechnic Institute.

Forrester Johnson, Aerodynamics Research, Boeing Commercial Airplane Group.

Robert W. MacCormack, Professor, Department of Aeronautics and Astronautics, Stanford University.

Stanley G. Rubin, Professor, Department of Aerospace Engineering and Engineering Mechanics, University of Cincinnati.

Manuel D. Salas, Director, Institute for Computer Applications in Science and Engineering, NASA Langley Research Center.

III. RESEARCH FELLOWS

Dimitri Mavriplis - Ph.D., Mechanical and Aerospace Engineering, Princeton University, 1988. Applied & Numerical Mathematics [Grid Techniques for Computational Fluid Dynamics]. (February 1997 to August 2001)

Piyush Mehrotra - Ph.D., Computer Science, University of Virginia, 1982. Computer Science [Programming Languages for Multiprocessor Systems]. (January 1991 to September 2004)

IV. SENIOR STAFF SCIENTISTS

Thomas W. Crockett - B.S., Mathematics, The College of William & Mary, 1977. Computer Science [System Software for Parallel Computing, Computer Graphics, and Scientific Visualization]. (February 1987 to October 1999)

R. Michael Lewis - Ph.D., Mathematical Sciences, Rice University, 1989. Applied & Numerical Mathematics [Multidisciplinary Design Optimization]. (May 1995 to August 2000)

Josip Lončarić - Ph.D., Applied Mathematics, Harvard University, 1985. Applied & Numerical Mathematics [Multidisciplinary Design Optimization]. (March 1996 to August 2001)

Li-Shi Luo - Ph.D., Physics, Georgia Institute of Technology, 1993. Computer Science [Parallel Algorithms]. (November 1996 to August 2002)

Zoubeida Ounaies - Ph.D., Engineering Science and Mechanics, The Pennsylvania State University, 1995. Physical Sciences [Characterization of Advanced Piezoelectric Materials]. (March 1999 to November 2002)

Alexander Povitsky - Ph.D., Mechanical Engineering, Moscow Institute of Steel and Alloys Technology (MISA), Russia, 1988. Computer Science [Parallelization and Formulation of Higher Order Schemes for Aeroacoustics Noise Propagation]. (October 1997 to August 2001)

Robert Rubinstein - Ph.D., Mathematics, Massachusetts Institute of Technology, 1972. Fluid Mechanics [Turbulence Modeling, Turbulence Management, and Acoustics]. (May 1998 to August 2002)

David Sidilkover - Ph.D., Applied Mathematics, The Weizmann Institute of Science, 1989. Applied & Numerical Mathematics [Numerical Analysis and Algorithms]. (November 1994 to May 2000)

Ji Su - Ph.D., Polymer Materials Science, Rutgers University, 1995. Physical Science [Electroactive Materials]. (January 2000 to November 2000)

V. SCIENTIFIC STAFF

Brian G. Allan - Ph.D., Mechanical Engineering, University of California at Berkeley, 1996. Applied & Numerical Mathematics [Multidisciplinary Design Optimization]. (February 1996 to November 2001)

Roger C. Hart - Ph.D., Physics, University of Tennessee, 1991. Physical Sciences [Measurement Science and Technology]. (December 1998 to October 2000)

Luc Huyse - Ph.D., Civil Engineering, Structures, University of Calgary, Canada, 1999. Applied & Numerical Mathematics [Managing Uncertainties in Multidisciplinary Research]. (October 1999 to August 2001)

Gerald Lüttgen - Ph.D., Computer Science, University of Passau, Germany, 1998. Computer Science [Formal Methods Research for Safety Critical Systems]. (October 1998 to August 2001)

César A. Muñoz - Ph.D., Computer Science, University of Paris 7, 1997. Computer Science [Formal Methods Research for Safety Critical Systems]. (May 1999 to April 2001)

VI. VISITING SCIENTISTS

Boris Diskin - Ph.D., Applied Mathematics, The Weizmann Institute of Science, Israel, 1998. Teaching Assistant, The Weizmann Institute of Science, Israel. Applied & Numerical Mathematics [Convergence Acceleration]. (July 1998 to September 2000)

Chaung Lin - Ph.D., Nuclear Engineering, University of California-Berkeley, 1983. Professor, Department of Engineering and System Science, National Tsing Hua University, Taiwan. Physical Sciences [Flow Modeling from Experimental Data and Real-time Control Using Predictive Control Techniques]. (August 1999 to January 2000)

Yeon-Gon Mo - Ph.D., Electrical Engineering, University of Nebraska at Lincoln, 1999. Graduate Research Assistant, Department of Electrical Engineering, University of Nebraska at Lincoln. Physical Science [Smart Materials and Flow Control]. (October 1999 to November 2000)

Sun Mok Paik - Ph.D., Physics, University of Maryland, 1988. Assistant Associate Professor, Department of Physics, Kangwon National University, Korea. Physical Science [Computational Materials]. (February 2000 to January 2001)

Linda Stals - Ph.D., Mathematics, Australian National University, 1996. Post-Doc, Department of Computer Science, Old Dominion University. Computer Science [Parallel Implicit Multilevel Algorithms]. (November 1998 to October 2000)

VII. SHORT-TERM VISITING SCIENTISTS

Pierre Lallemand - Ph.D., Physics, Universite de Paris, 1966. Director of Research, Centre National de la Recherche Scientifique, A.S.C.I., Universite Paris-Sud. Computer Science. (March 2000 to April 2000)

VIII. ASSOCIATE RESEARCH FELLOW

David E. Keyes - Ph.D., Applied Mathematics, Harvard University, 1984. Computer Science [Parallel Numerical Algorithms]

IX. CONSULTANTS

Ponnampalam Balakumar - Ph.D., Aeronautics and Astronautics, Massachusetts Institute of Technology, 1986. Associate Professor, Department of Aerospace Engineering, Old Dominion University. Fluid Mechanics [Stability and Transition]

H. Thomas Banks - Ph.D., Applied Mathematics, Purdue University, 1967. Professor, Department of Mathematics, Center for Research in Scientific Computations, North Carolina State University. Applied & Numerical Mathematics [Control Theory]

Richard W. Barnwell - Ph.D., Engineering Mechanics, Virginia Polytechnic Institute and State University, 1968. Professor, Department of Aerospace and Ocean Engineering, Engineering Science and Mechanics. Virginia Polytechnic Institute and State University. Fluid Mechanics [Turbulence Modeling]

Oktay Baysal - Ph.D., Mechanical Engineering, Louisiana State University, 1982. Eminent Scholar and Professor, Department of Aerospace Engineering, Old Dominion University. Applied & Numerical Mathematics

Achi Brandt - Ph.D., Mathematics, The Weizmann Institute of Science, 1965. Professor, Department of Applied Mathematics, The Weizmann Institute of Science, Israel. Applied & Numerical Mathematics [Convergence Acceleration]

Thomas W. Crockett - B.S., Mathematics, The College of William & Mary, 1977. Senior Research Associate, Computational Science Cluster, The College of William & Mary. Computer Science [Scientific Visualization]

Ayodeji O. Demuren - Ph.D., Mechanical Engineering, Imperial College London, United Kingdom, 1979. Associate Professor, Department of Mechanical Engineering and Mechanics, Old Dominion University. Fluid Mechanics [Numerical Modeling of Turbulent Flows]

Earl H. Dowell - Ph.D., Aeronautics & Astronautics, Massachusetts Institute of Technology, 1964. J.A. Jones Professor, Department of Mechanical Engineering, Duke University. Physical Sciences [Autonomous System Identification and Control]

David Gottlieb - Ph.D., Numerical Analysis, Tel-Aviv University, Israel, 1972. Ford Foundation Professor & Chair, Division of Applied Mathematics, Brown University. Applied & Numerical Mathematics [Boundary Conditions for Hyperbolic Systems]

Jan S. Hesthaven - Ph.D., Applied Mathematics/Numerical Analysis, Technical University of Denmark, 1995. Visiting Assistant Professor, Division of Applied Mathematics, Brown University. Physical Sciences [Computational Electromagnetics]

Fang Q. Hu - Ph.D., Applied Mathematics, Florida State University, 1990. Assistant Professor, Department of Mathematics and Statistics, Old Dominion University. Fluid Mechanics [Aeroacoustics]

Osama A. Kandil - Ph.D., Engineering Mechanics, Virginia Polytechnic Institute and State University, 1974. Professor and Eminent Scholar of Aerospace Engineering and Chair, Department of Aerospace Engineering, Old Dominion University. Applied & Numerical Analysis [Computational Fluid Dynamics]

Frank Kozusko - Ph.D., Computational and Applied Mathematics, Old Dominion University, 1995. Assistant Professor, Department of Mathematics, Hampton University. Fluid Mechanics [Airfoil Design]

David G. Lasseigne - Ph.D., Applied Mathematics, Northwestern University, 1985. Assistant Professor, Department of Mathematics and Statistics, Old Dominion University. Fluid Mechanics [Asymptotic and Numerical Methods for Computational Fluid Dynamics]

James E. Martin - Ph.D., Applied Mathematics, Brown University, 1991. Assistant Professor, Department of Mathematics, Christopher Newport University. Fluid Mechanics [Turbulence and Computation]

Jan Nördstrom - Ph.D., Numerical Analysis, Uppsala University, Sweden, 1993. Senior Scientist, The Aeronautical Research Institute of Sweden. Applied & Numerical Mathematics [Global Boundary Conditions for Aerodynamic and Aeroacoustic Computations]

Alex Pothen - Ph.D., Applied Mathematics, Cornell University, 1984. Professor, Department of Computer Science, Old Dominion University. Computer Science [Parallel Numerical Algorithms]

Viktor Ryaben'kii - Ph.D., Stability of Difference Equations, Moscow State University, 1953. Leading Research Scientist, Keldysh Institute for Applied Mathematics, Russian Academy of Sciences and Full Professor, Department of Control and Applied Mathematics, Moscow Institute of Physics and Technology. Applied & Numerical Mathematics [Global Boundary Conditions for Aerodynamic and Aeroacoustic Computations]

Chi-Wang Shu - Ph.D., Mathematics, University of California-Los Angeles, 1986. Associate Professor, Division of Applied Mathematics, Brown University. Fluid Mechanics [Computational Aeroacoustics]

Ralph C. Smith - Ph.D., Mathematics, Montana State University, 1990. Assistant Professor, Department of Mathematics, Iowa State University. Applied & Numerical Mathematics [Optimal Control Techniques for Structural Acoustics Problems]

Virginia Torczon - Ph.D., Mathematical Sciences, Rice University, 1989. Assistant Professor, Department of Computer Science, The College of William & Mary. Computer Science [Parallel Algorithms for Optimization Including Multidisciplinary Optimization]

Michael W. Trosset - Ph.D., Statistics, University of California-Berkeley, 1983. Department of Mathematics, The College of William & Mary. Applied & Numerical Mathematics [Multidisciplinary Optimization]

Semyon V. Tsynkov - Ph.D., Computational Mathematics, Keldysh Institute for Applied Mathematics, Russian Academy of Sciences, 1991. Senior Lecturer, Department of Applied Mathematics, Tel Aviv University, Israel. Applied & Numerical Mathematics [Global Boundary Conditions for Aerodynamic and Aeroacoustic Computations]

Xiaodong Zhang - Ph.D., Computer Sciences, University of Colorado-Boulder, 1989. Professor, Department of Computer Science, The College of William & Mary. Computer Science

Mohammad Zubair - Ph.D., Computer Science, Indian Institute of Technology, Delhi, India, 1987. Professor, Department of Computer Science, Old Dominion University. Computer Science [Performance of Unstructured Flow-solvers on Multi-processor Machines]

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503.

1. AGENCY USE ONLY *(Leave blank)*	2. REPORT DATE	3. REPORT TYPE AND DATES COVERED
	April 2000	Contractor Report

4. TITLE AND SUBTITLE
Semiannual Report
October 1, 1999 through March 31, 2000

5. FUNDING NUMBERS
C NAS1-97046
WU 505-90-52-01

6. AUTHOR(S)

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)
Institute for Computer Applications in Science and Engineering
Mail Stop 132C, NASA Langley Research Center
Hampton, VA 23681-2199

8. PERFORMING ORGANIZATION REPORT NUMBER

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)
National Aeronautics and Space Administration
Langley Research Center
Hampton, VA 23681-2199

10. SPONSORING/MONITORING AGENCY REPORT NUMBER
NASA/CR-2000-210105

11. SUPPLEMENTARY NOTES
Langley Technical Monitor: Dennis M. Bushnell
Final Report

12a. DISTRIBUTION/AVAILABILITY STATEMENT
Unclassified–Unlimited
Subject Category 59
Distribution: Nonstandard
Availability: NASA-CASI (301) 621-0390

12b. DISTRIBUTION CODE

13. ABSTRACT *(Maximum 200 words)*
This report summarizes research conducted at the Institute for Computer Applications in Science and Engineering in applied mathematics, computer science, fluid mechanics, and structures and materials during the period October 1, 1999 through March 31, 2000.

14. SUBJECT TERMS
applied mathematics, multidiscipinary design optimization, fluid mechanics, turbulence, flow control, acoustics, materials, structures, computer science, system software, systems engineering, parallel algorithms

15. NUMBER OF PAGES
49

16. PRICE CODE
A04

17. SECURITY CLASSIFICATION OF REPORT	18. SECURITY CLASSIFICATION OF THIS PAGE	19. SECURITY CLASSIFICATION OF ABSTRACT	20. LIMITATION OF ABSTRACT
Unclassified	Unclassified		